DREO CR 1999-116

ANALYSIS OF REAL TIME RESULTS FROM DIINS' 1995 SEA TRIALS (U)

by

D.A.G. Arden

DALE ARDEN CONSULTING

REPORT NO. DAC-01-97

August 1997

Ottawa

ANALYSIS OF REAL TIME RESULTS FROM DIINS' 1995 SEA TRIALS (U)

by
D.A.G. Arden

DALE ARDEN CONSULTING

REPORT NO. DAC-01-97

August 1997
Ottawa

EXECUTIVE SUMMARY

The Dual Inertial Integrated Navigation System (DIINS) is a Kalman filter based, multi-sensor, optimally integrated navigation system developed at the Defence Research Establishment Ottawa. Development of DIINS was initiated with a view to satisfying the navigational requirements of the Canadian Navy's TRUMP, CPF and future submarines. DIINS has been designed to provide a single source of navigation information (position, velocity, attitude etc.) to other ship board systems with minimize data latency, improve navigation reliability and accuracy, and provide automatic sensor failure detection.

DIINS differs from earlier integrated navigation systems developed at DREO in two main areas: it is based on (two) inertial navigation systems; and it was designed to have a number of filters running concurrently.

The specific architecture selected for DIINS has been named a Hierarchic Alliance of Filters. It consists of several independently operating Kalman filters arranged in a "root" structure. Each of two main branches is based on data from one INS and consists of a number of filter levels. At the top is the "unifilter", which uses all available aiding sensors. Filters in subsequent levels exclude various combinations of aiding sensors in an arrangement that ensures there is always at least one uncorrupted filter, regardless of the sensor failure status.

The highest level filter which is safe from corruption by a sensor failure is eligible for use as DIINS' output source. In most circumstances, one of the unifilters is used. If an aiding sensor failure occurs which has the potential to corrupt a unifilter, the best filter in the next lower branch is selected as the output source.

Each year from 1992 to 1995, DIINS sensors were mounted onboard the research ship, CFAV Endeavour, based in Esquimalt, B.C. Sensor data was collected and real time tests were performed over a period of 1–2 weeks. This report describes latest DIINS Sea Trial which covered the time period from May 2, 1995 (Julian Day 123) to May 15, 1995 (Julian Day 136). Onboard the Endeavour for these trials were two MK–29 MOD 3 INSs built by Sperry Marine Inc., a Trimble Navigation P(Y) code military GPS receiver and an Internav LC360 Loran-C receiver. These are main navigation sensors used on the CPFs. Also, onboard was a differential GPS system operated by the University of Calgary. This sensor acted as the position reference system. It has also been used as a DIINS aiding sensor in post-mission tests.

The report begins with a description of all the navigation sensors involved in ST95. This is followed by a description of the real-time DIINS filter results. The report concludes with an analysis of the real-time FDIR results.

Contents

EXECUTIVE SUMMARY ii

List of Tables vii

List of Figures viii

Chapter 1 INTRODUCTION 1

1.1 A Brief Description of DIINS 1

1.2 Summary of the 1995 DIINS Sea Trials 2

1.3 Summary of This Report 6

Chapter 2 POSITION REFERENCE: DESCRIPTION AND ANALYSIS 7

2.1 Position Reference Description 7

2.2 Position Reference Analysis 8

2.2.1 Addendum 13

Chapter 3 NAVIGATION SENSORS: DESCRIPTION AND ANALYSIS 15

3.1 Trimble TASMAN P(Y) Code GPS Receiver 15

3.1.1 Data Formats 15

3.1.2 Position Errors vs Differential GPS 20

3.1.3 Filter Models 22

3.2 Sperry MK29 Inertial Navigation Systems 28

3.2.1 Position Errors 28

3.2.2 SKOR Manual Fix Accuracy 32

3.2.3 Status Messages 45

3.2.4 Filter Models 49

3.3 Internav LC360 Loran-C Receiver 50

3.3.1 Time Difference Errors 50

3.3.2 Filter Models 55

Chapter 4	REAL-TIME FILTER RESULTS: DESCRIPTION AND ANALYSIS	57
4.1	Filter 1 (INS 1 with INS 2 Position Measurements)	57
4.2	Filter 2 (INS 1 with TASMAN P(Y) GPS Position and LC360 Loran-C Time Difference Measurements)	58
4.2.1	Anomaly 1	58
4.2.2	Anomaly 2	61
4.3	Filter 3 (INS 2 with TASMAN P(Y) GPS Positions and LC360 Loran-C Time Differences)	66
4.3.1	Anomaly 1	67
4.4	Filter 4 (INS 1 with TASMAN P(Y) GPS Positions)	79
4.5	Filter 5 (INS 2 with TASMAN P(Y) GPS Positions)	79
4.6	Filter 6 (INS 1 with LC360 Loran-C Time Differences)	79
4.6.1	Anomaly 1	80
4.7	Filter 7 (INS 2 with LC360 Loran-C Time Differences)	82
Chapter 5	REAL TIME FAILURE DETECTION, ISOLATION AND RECONFIGURATION	83
5.1	INS Status Messages	83
5.1.1	Run 1, Days 129–132	86
5.1.2	Run 2, Day 133	88
5.1.3	Run 3, Days 134–135	89
5.1.4	Run 4, Day 136	90
5.2	Measurement Analysis (Hard Failures or Data Problems)	91
5.2.1	Data Problems	92
5.2.2	Hard Failures	93
5.3	Chi-Squared Tests (Soft Failures)	96
5.4	Selection of the Best Filter for Output	99

Chapter 6 SUMMARY 101
References 106
Appendix A Post-Mission Position Fix Data Summaries . . . 107
A.1 MK29–1 Fix Summary 108
A.1.1 Run 1, Days 129–132 108
A.1.2 Run 2, Day 133 108
A.1.3 Run 3, Days 134–135 109
A.2 MK29–2 Fix Summary 109
A.2.1 Run 1, Days 129–132 109
A.2.2 Run 2, Day 133 110
A.2.3 Run 3, Day 134–135 110

List of Tables

Table 1	DIINS Sea Trial 1995 Navigation Sensor Body Frame Coordinates	3
Table 2	Description of U of C Differential GPS Data Sets	8
Table 3	Description of TASMAN P(Y) GPS Data Sets	9
Table 4	Analysis of Gaps in U of C Differential GPS Data	12
Table 5	dGPS — P(Y) Position Differences	20
Table 6	dGPS — P(Y) Position Differences	22
Table 7	Stationary P(Y) Position Errors	23
Table 8	P(Y) Position Errors vs dGPS, ST94	23
Table 9	P(Y) Position Errors vs dGPS, ST94 — Gauss-Markov Model	24
Table 10	P(Y) Position Errors vs dGPS, ST95 — Gauss-Markov Model	24
Table 11	P(Y) Position Errors vs dGPS, Van Tests August 1994	25
Table 12	P(Y) Position Errors vs dGPS, Van Tests 1994 — Gauss-Markov Model	25
Table 13	ST95 INS Position Reset Schedule	31
Table 14	Radial Position Errors Before and After All ST95 Fixes	34
Table 16	Post-Mission Residual Test Statistics	40
Table 15	Errors in Fixes Entered for ST95	41
Table 17	Loran-C TD Error Models from DIINS ST94 and ST95	56
Table 18	Soft Failure Isolation	97

List of Figures

Figure 1	DIINS Sea Trial 1995 Navigation Sensor Locations	3
Figure 2	dGPS Data Gaps, Days 129–132	10
Figure 3	dGPS Data Gaps, Days 133	10
Figure 4	dGPS Data Gaps, Days 134–135	11
Figure 5	dGPS Data Gaps, Days 133	11
Figure 6	Monitor to Ship Distances	13
Figure 7	TASMAN Latitude Errors vs dGPS with TASMAN Predicted Bounds, Days 131–132	17
Figure 8	TASMAN Longitude Errors vs dGPS with TASMAN Predicted Bounds, Days 131–132	17
Figure 9	TASMAN Height Errors vs dGPS with TASMAN Predicted Bounds, Days 131–132	18
Figure 10	TASMAN Latitude Errors vs dGPS with Expected Bounds, Days 135–136	18
Figure 11	TASMAN Longitude Errors vs dGPS with Expected Bounds, Days 135–136	19
Figure 12	TASMAN Height Errors vs dGPS with TASMAN Predicted Bounds, Days 135–136	19
Figure 13	TASMAN P(Y) Latitude Differences vs dGPS	21
Figure 14	TASMAN P(Y) Longitude Differences vs dGPS	21
Figure 15	TASMAN P(Y) Height Differences vs dGPS	22
Figure 16	MK29–1 Latitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds	29
Figure 17	MK29–1 Longitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds	29
Figure 18	MK29–2 Latitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds	30
Figure 19	MK29–2 Longitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds	30
Figure 20	MK29 Position Errors with Velocity Damping On and Off	32
Figure 21	MK29–2 Latitude Errors Across a Reset with SKOR Predicted Bounds	35

Figure 22	MK29–2 Longitude Errors Across a Reset with SKOR Predicted Bounds	35
Figure 23	MK29–1 Performance Degraded Status Flag	47
Figure 24	MK29–1 Inertial Well Settled Flag	47
Figure 25	MK29–2 Performance Degraded Status Flag	48
Figure 26	MK29–2 Inertial Well Settled Flag	48
Figure 27	Loran-C Y and Z TD Errors, ST95, Days 129–132	51
Figure 28	Loran-C Y and Z TD Errors, ST94, Days 322–325	51
Figure 29	DIINS Sea Trials' Trajectory	53
Figure 30	3–D View of Loran-C Y TD Errors, ST94, Days 321–325	54
Figure 31	3–D View of Loran-C Y TD Errors, ST95, Days 129–132	54
Figure 32	Filter2: Position Error Anomalies	59
Figure 33	Filter2: Position Standard Deviation Anomaly	62
Figure 34	Filter 2 Latitude Residual with Failure Bounds (5 sigma)	63
Figure 35	Filter 2 Latitude Residual — Before Anomaly	64
Figure 36	Filter 2 Latitude Residual — After Anomaly	64
Figure 37	Filter 3 Latitude Errors with DIINS Standard Deviations, Day 134	69
Figure 38	Filter 3 Longitude Errors with DIINS Standard Deviations, Day 134	69
Figure 39	MK29–2 Latitude Error vs GPS with SKOR Estimated Standard Deviations, Day 133-134	70
Figure 40	MK29–2 Longitude Error vs GPS with SKOR Estimated Standard Deviations, Day 133-134	70
Figure 41	MK29–2 Inertial Velocity Damping ON Flag, Day 133-134	73
Figure 42	MK29–2 Velocity Damping Mode Flag, Day 133-134	73
Figure 43	DIINS Trajectory, Day 133-134	74
Figure 44	DIINS Longitude, Day 133-134	74
Figure 45	MK29–2 East Velocity Error vs East Ocean Current Estimate	75

Figure 46 MK29–2 North and East Velocity Corrections (with Damp Flag), Day 133 75
Figure 47 MK29–2 Performance Degraded Status Flag, Day 133–134 76
Figure 48 MK29–2 Inertial Well Settled Status Flag, Day 133-134 76

Chapter 1: INTRODUCTION

1.1: A Brief Description of DIINS

The Dual Inertial Integrated Navigation System (DIINS) is a Kalman filter based, multi-sensor, optimally integrated navigation system developed at the Defence Research Establishment Ottawa. Development of DIINS was initiated with a view to satisfying the navigational requirements of the Canadian Navy's TRUMP, CPF and future submarines. The primary objectives of DIINS are to:

1. provide single source of navigation information (position, velocity, attitude etc.),
2. improve navigation reliability,
3. improve navigation accuracy,
4. minimize data latency to other ship board systems,
5. provide automatic sensor failure detection.

all within the constraints of the available sensor suite.

DIINS follows on earlier work at DREO:

1. MINS (Marine Integrated Navigation System) is the first generation marine optimally integrated navigation system developed at DREO. It was based on a gyrocompass, speedlog and the radionavigation sensors available when MINS was developed. MINS was installed on many Navy ship.
2. PLANS (Primary Land Arctic Navigation System) is an optimally integrated navigation system designed and developed at DREO to work onboard a land vehicle at high latitudes.

DIINS differs from these earlier systems in two main areas:

1. it is based on (two) inertial navigation systems,
2. and it was designed to have a number of filters running concurrently.

The first change required the analysis and development of fairly complex inertial error models. The second change was imposed by the requirement to maximize the reliability of the navigation system. It was also aided by the extra redundancy supplied by having two INSs available.

The specific architecture selected for DIINS is being called a Hierarchic Alliance of Filters. It consists of several independently operating Kalman filters

arranged in a "root" structure. At the top of the structure is a filter which combines data from each of the INSs. This filter is used exclusively for failure isolation of the INSs. Directly below this filter are two "unifilters", each based on one INS and using measurements from all available aiding sensors. Branching out from each unifilter are n additional filters (where n is the number of aiding sensors), each of which excludes measurements from one aid. The next lower level of branching is composed of the set of filters which each exclude two aiding sensors. The lowest level of branching contains n filters which each use one of aids. With this architecture, there are n levels (in two branches) below the INS1/INS2 filter.

In most circumstances, output from one of the unifilters is used for navigation output. If an aiding sensor failure occurs which has the potential to corrupt a unifilter, the filter in the next lower branch which does not use the failed sensor is selected as the new output source. Failure isolation of an INS removes that INS' branch of filters from consideration for output: the best filter in the healthy INS' branch would be selected. The corruptible failure of a second aiding sensor would lead to the filter in the next lower level which excludes both failed sensors. If the second INS fails, unfiltered data from the best aiding sensor would be used for output. Recovery of a failed sensor can result in the reversal of this output selection procedure. This architecture provides an optimal uncorrupted filter under most failure conditions, and unfiltered aiding sensor output if both INSs have failed.

Details of the DIINS Kalman filter algorithms and error models can be found in reference [1]. The DIINS FDIR processor is described in reference [2].

1.2: Summary of the 1995 DIINS Sea Trials

Each year from 1992 to 1995, DIINS sensors were mounted onboard the research ship, CFAV Endeavour, based in Esquimalt, B.C. Sensor data was collected and real time tests were performed over a period of 1–2 weeks. The latest and last (?) DIINS Sea Trial covered the period from May 2, 1995 (Julian Day 123) to May 15, 1995 (Julian Day 136). Onboard the Endeavour for these trials were two MK–29 MOD 3 INSs built by Sperry Marine Inc., a Trimble Navigation P(Y) code military GPS receiver and an Internav LC360 Loran-C receiver. These are main navigation sensors used on the CPFs and the sensors integrated by DIINS. Also, onboard as the position reference system was a differential GPS system operated by the University of Calgary. Figure 1 shows the approximate relative

locations of all sensors. Table 1 gives the coordinates of all sensors relative to the forward MK29(-1) in body frame coordinates (positive forward, positive starboard, positive down). Note that all real-time DIINS output is referred to the

Figure 1 DIINS Sea Trial 1995 Navigation Sensor Locations

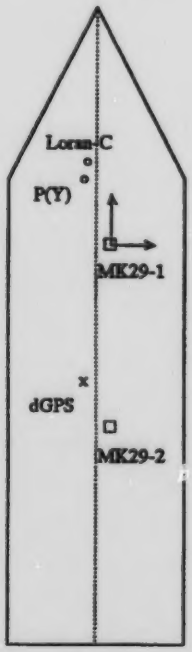

Table 1 DIINS Sea Trial 1995 Navigation Sensor Body Frame Coordinates

Sensor	x^b Coordinate (m)	y^b Coordinate (m)	z^b Coordinate (m)
MK29-1	0.00	0.00	0.00
MK29-2	-22.72	0.00	0.00
P(Y) GPS	6.53	-2.35	-6.81
Loran-C	7.75	-2.55	-6.80
dGPS	-16.00	-1.36	-9.46
DIINS	0.00	0.00	0.00

MK29–1 position.

In the ST95 configuration, DIINS runs seven Kalman filters concurrently:

1. One filter using one INS as the aiding sensor for the other. This filter is used only for failure analysis.
2. There are two filters (one based on each INS) which use GPS position measurements and Loran-C time difference measurements. These are the unifilters which are used for final DIINS output under most circumstances.
3. There are two filters (one based on each INS) which use only GPS position measurements.
4. There are two filters (one based on each INS) which use only Loran-C measurements.

The last four filters in the list comprise the second branches of Hierarchic Alliance architecture: they are used for output under certain sensor failure conditions. Reference [2] describes the sensor Failure Detection, Isolation and Reconfiguration concepts (including selection of the output source) in some detail.

Useful test data was collected in four separate runs during the 1995 sea trial (ST95). A "run" is defined as the time interval from DIINS' initialization to termination. Termination can result from a graceful, planned system turn off, or it can result from an unexpected software or hardware "crash". The four runs identified for ST95 indicate that DIINS was turned off (or crashed) four times during these trials. A brief summary of these runs follows.

1. Run 1, Days 129–132:
   a. 74,276,507.78 seconds (129:16:22) — INSs taken out of Dockside, ship under way.
   b. Departure to 74,315,378.81 seconds (130:03:10) — INS resets every two hours.
   c. 74,344,621.68 seconds (130:11:17) and 74,451,842.14 seconds (131:17:04) — INS resets.
   d. 132:03:00 — tied up at fuel jetty.
   e. 74,490,903.89 seconds (132:03:55) — INSs put back in Dockside. DIINS halted by operator.
2. Run 2, Day 133:
   a. 74,622,714.42 seconds (133:16:32) — INSs taken out of Dockside.

b. 133:16:59 — ship under way.
c. Departure to 74,657,530.41 seconds (134:02:12) — INS resets every two hours.
d. 134:05:21 — DIINS crashed, cause unknown.

3. Run 3, Days 134–135:

a. 74,674,725.93 seconds (134:06:59) — DIINS restarted at sea.
b. 74,686,624.81 seconds (134:10:17) and 74,757,233.43 seconds (135:05:54) — two sparse INS resets.
c. 74,822,546.69 (136:00:02) — DIINS crashed because of a Loran-C receiver malfunction.

4. Day 136:

a. 74,823,808.77 seconds (136:00:23) — DIINS restarted at sea. There were no INS resets given over the last 34 hours.
b. 74,881,067.00 seconds (136:16:18) — end of trials. INSs put back in Dockside. DIINS halted by operator.

Note: exact times refer to DIINS log messages for INS1 from Filter 1 (INS1/INS2). Times are given in "DIINS" time (defined below) and the equivalent Julian day number:hours:minutes. The conversion from DIINS time to Julian time is also described below.

The DIINS time used during ST95 has its origin at January 1, 1993, 00:00:00 and has units of seconds. DIINS time can be converted to yyyy:dd.dddd and yyyy:dd:hh:mm:ss.ss format as follows:

1. Convert seconds elapsed to days elapsed

$$dd.ddd, ddd, d_{93} = ss, sss, sss.ss_{DIINS} \frac{1 \text{ day}}{86,400 \text{ s/day}} \quad (1)$$

2. Convert time elapsed (in days) since Jan01, 1993 to time elapsed (in days) since Jan01 of the current year, eg. 1995, to get the Julian Day

$$\begin{aligned} dd.ddd, ddd, d_{95} &= dd.ddd, ddd, d_{93} - N * 365 \\ &= dd.ddd, ddd, d_{93} - 2 * 365 \\ &= dd.ddd, ddd, d_{93} - 730 \end{aligned} \quad (2)$$

3. Check the current year

$$\begin{aligned} yyyy_{1995} &= 1993 + N \\ &= 1993 + 2 \\ &= 1995 \end{aligned} \tag{3}$$

4. Convert the fractional days to hh:mm:ss.ss

a.

$$hh.hhhh_{dd} = 0.ddd, ddd, d_{95} \frac{24 \text{ hours}}{1 \text{ day}} \tag{4}$$

b.

$$mm.mmm_{hh} = hh.hhhh_{dd} \frac{60 \text{ min}}{1 \text{ hour}} \tag{5}$$

c.

$$ss.ss_{mm} = mm.mmm_{hh} \frac{60 \text{ sec}}{1 \text{ min}} \tag{6}$$

1.3: Summary of This Report

This report describes all aspects of the DIINS 1995 Sea trial. It begins with a description of the differential GPS position reference system. This is followed by descriptions of all the DIINS sensors: the Sperry MK29 Mod 3 INSs, the Trimble TASMAN P(Y) code GPS receiver and the Internav LC-360 Loran-C receiver. Next is a description of the real-time DIINS filter results. The report concludes with an analysis of the real-time FDIR results.

Chapter 2: POSITION REFERENCE: DESCRIPTION AND ANALYSIS

The DIINS version used for ST95 output and saved the data from each navigation sensor as well as the navigation results produced by every filter. An independent position reference system was also onboard specifically to help in the analysis of the results. This chapter will describe the reference system and analyze its performance.

2.1: Position Reference Description

Differential GPS position reference data was provided by the University of Calgary (U of C). Details of the procedures can be found in reference [3]; a brief summary follows.

The U of C dGPS system consisted of two 3951 NovAtel GPSCard™ receivers. Each came equipped with a NovAtel 501 antenna mounted in a choke ring ground plane and other necessary accessory hardware and software. The mobile antenna was mounted on the roll axis of the ship, to the stern of the pitch axis and about 9 metres above the main deck. In body frame coordinates (forward, starboard, down), with origin at the DIINS output reference point (coincident with the forward INS1), the coordinates of the antenna in metres are (–16.0, –1.36, –9.46). The stationary reference antenna was mounted on a 30 foot tower on the roof the Defence Research Establishment Pacific (DREP) building. Raw data was collected and stored using laptop computers. About once per day, data had to be dumped from the laptops to optical storage devices. All dGPS processing was done post mission at the U of C using their C^3NAV^{TM} software: carrier smoothed code measurements from the monitor are used to differentially correct carrier smoothed code measurements from the ship's receiver; then a least squares adjustment is used to compute positions on an epoch by epoch basis.

The final dGPS positions were sent to DREO in ASCII format via computer diskette and transferred to the DREO Navigation Group's Sun workstation network. Position data records each contain GPS time of week (in seconds), latitude and longitude (in degrees and decimal degrees), and ellipsoid height (in metres). The U of C ASCII files were then converted binary format: the data is the same except that the time is converted to DIINS format in the binary files. NOTE: since

the U of C data does not include GPS week numbers, the conversion program is only valid for Julian days 127–140 inclusive.

2.2: Position Reference Analysis

The full U of C dGPS data set is described in the following table.

Table 2 Description of U of C Differential GPS Data Sets

Data Set (J-days)	128-132	133-135
Start Time, seconds (ddd:hh:mm)	74,191,759 (128:16:49)	74,622,811 (133:16:33)
Stop Time, seconds (ddd:hh:mm)	74,487,461 (132:02:58)	74,878,168 (136:15:29)
Elapsed Time, seconds (d:hh:mm)	295,702 (3:10:08)	255,357 (2:22:56)
No. of Records	194,090	235,544

Notice the elapsed time versus number of records: at a nominal 1 second output rate, these numbers should be approximately equal. In the first data set (Julian Days 128–132), over 100,000 records, representing more than one in every three records, are missing. In the second data set (Julian Days 133–135), almost 20,00 records, representing about one in every thirteen records, are missing. Overall, 429,634 out of a possible 551,059 records (or 78%) are available. The missing records can be attributed to problems with either the reference or mobile station or to problems computing or applying the differential corrections (eg. different satellites being tracked by the two receivers).

The TASMAN operated as a stand-alone P(Y) code receiver and, therefore, provided data more reliably. However, data was collected from the TASMAN by the DIINS software: data was collected only when DIINS was running. Therefore,

the TASMAN data sets are different from the U of C data sets. The following table describes the TASMAN (and also other DIINS output) data sets.

Table 3 Description of TASMAN P(Y) GPS Data Sets

Data Set (J-days)	129-132	133	134-135	136
Start Time, seconds (ddd:hh:mm)	74,276,551 (129:16:23)	74,621,667 (133:16:14)	74,674,687 (134:06:58)	74,823,769 (136:00:23)
Stop Time, seconds (ddd:hh:mm)	74,491,097 (132:03:58)	74,668,910 (134:05:22)	74,822,560 (136:00:03)	74,881,069 (136:16:18)
Elapsed Time, seconds (d:hh:mm)	214,543 (2:11:36)	47,243 (0:13:07)	147,873 (1:17:05)	57,300 (0:15:55)
No. of Records	214,379	47,060	147,870	57,282

Note the high reliability of the TASMAN data: 466,591 of a possible 466,959 records were collected — an average of 99.9%.

To identify specific data gaps in the U of C dGPS data, it was compared to the TASMAN data. In fact, a position error program called PositionDiff.ada was used. This program reads in position data from 2 files and computes position differences. In doing so, it requires input of the maximum allowable time difference between the 2 records being compared. If the limit is exceeded, the difference is not computed, and the time difference is saved in a special file. The program was modified slightly to output 1 when the limit was not exceeded (ie. dGPS available) and 0 when the limit was exceeded (ie. dGPS NOT available). A summary file is also output which echoes all of the input parameters, gives the times of the first and last records successfully differenced, tabulates the statistical summary of all the differences, and provides counts of successful and rejected differences.

Figure 2 dGPS Data Gaps, Days 129–132

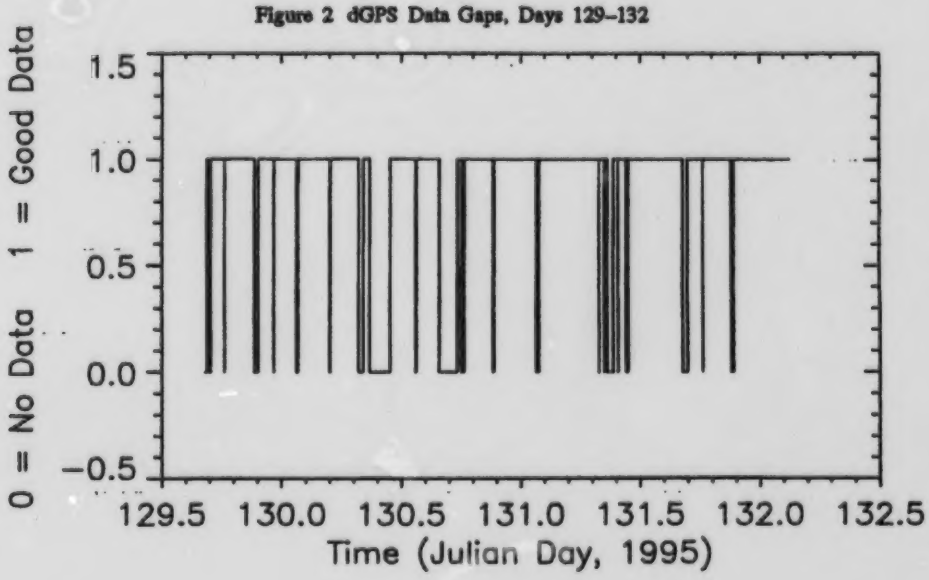

Figure 3 dGPS Data Gaps, Days 133

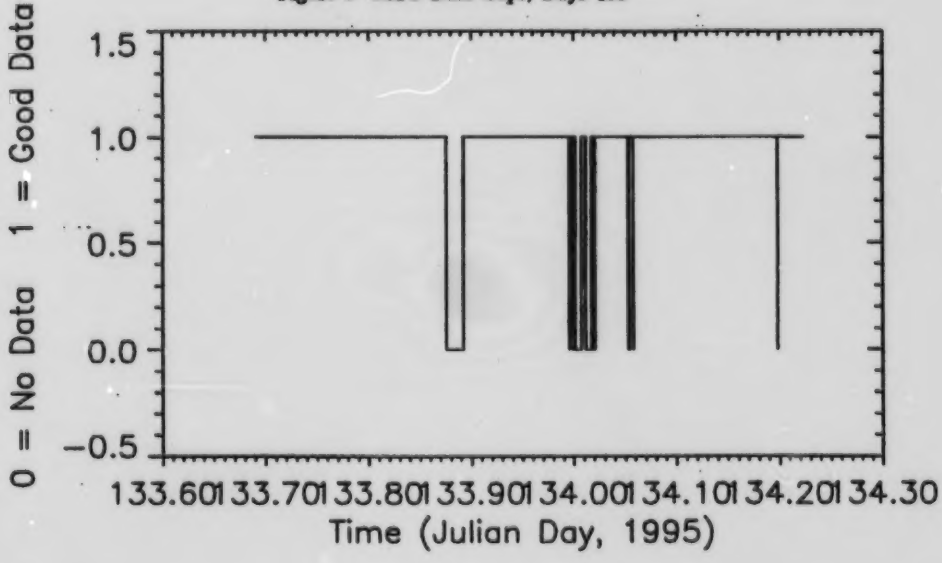

Figure 4 dGPS Data Gaps, Days 134–135

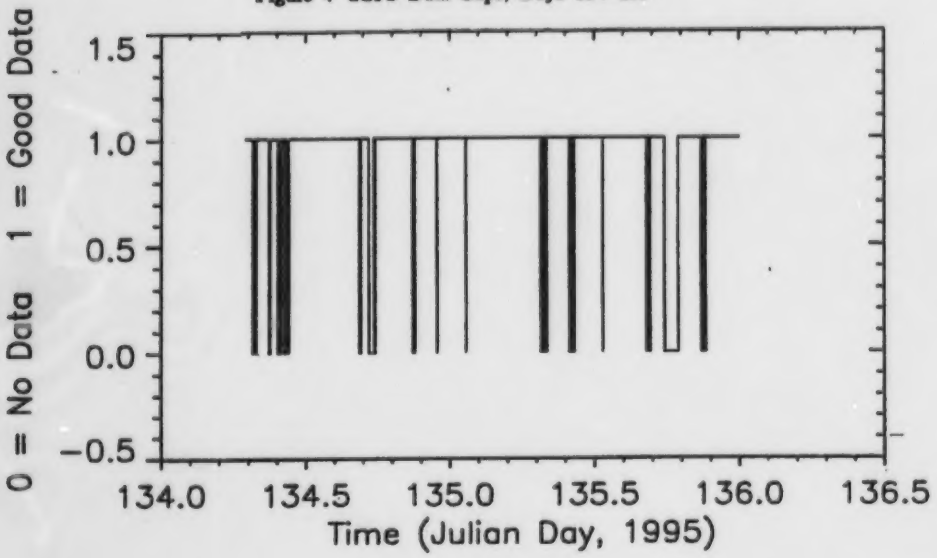

Figure 5 dGPS Data Gaps, Days 133

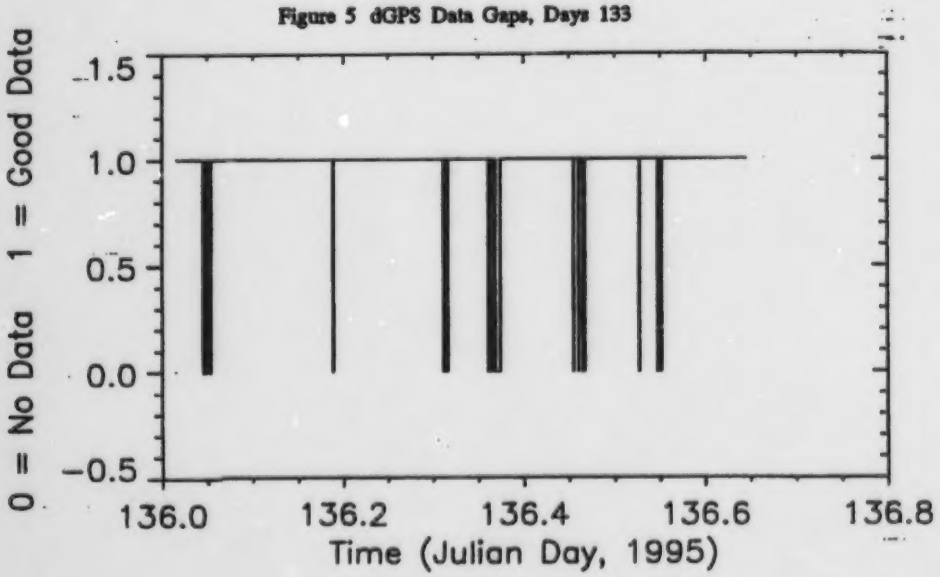

A time difference limit of 10 seconds was chosen in an attempt to eliminate the small data gaps which may be present in the TASMAN data. NOTE: since both receivers output data at nominally one second intervals, a time difference limit of 0.5 seconds will identify ALL data gaps. This would be the preferred value for use in position difference analysis. The following table summarizes the results of the time difference tests for each DIINS data set.

Table 4 Analysis of Gaps in U of C Differential GPS Data

Data Set (J-days)	129-132	133	134-135	136
No. Records Processed	185,351	43,172	138,727	52,760
No. Records Rejected (% of Total)	25,392 (12)	2,743 (6)	9,141 (6)	1,626 (3)

In total, 458,912 records were processed and 47,872 were rejected for an overall 8% rejection rate.

Figures 2 to 5 graphically illustrate the reliability of the dGPS data.

The accuracies expected of the dGPS results are not directly available. However, they can be inferred from available information. The University of Calgary supplied pseudorange residuals for all of the dGPS processing. There is one residual every second for every satellite used in the dGPS position calculations. In total, 2,146,963 residuals were computed. The mean and standard deviation for all residuals were computed:

- Mean = 0.0021 metres;
- Standard deviation = 0.405 metres.

All residuals were equally weighted. RMS would have been more appropriate, but only the standard deviation was available from existing software, and since the mean is very small, the standard deviation and RMS should be almost equal.

Further, analysis of the dGPS results from the 1994 DIINS sea trial were provided by the U of C (reference [4]). ST94 and ST95 had similar objectives, used similar equipment and were carried out in the same geographical area. The analysis used C^3NAV^{TM} and a more accurate procedure using carrier phase only

(U of C's program, FLYKIN™). FLYKIN™ results were expected to produce a trajectory accurate to about 10 centimetres. A total of about 4 hours of data, at monitor to ship distances of up to 80 kilometres, produced RMS differences of 0.27, 0.45 and 0.91 metres in latitude, longitude and height, respectively. At larger distances, FLYKIN™ loses its effectiveness. Reference [4] states that errors are expected to "grow slightly" up to about 200 kilometres Figure 6 shows monitor to ship distances for the complete 1995 trials. Distances are typically about 150 kilometres, with maximums of about 220 kilometres

Figure 6 Monitor to Ship Distances

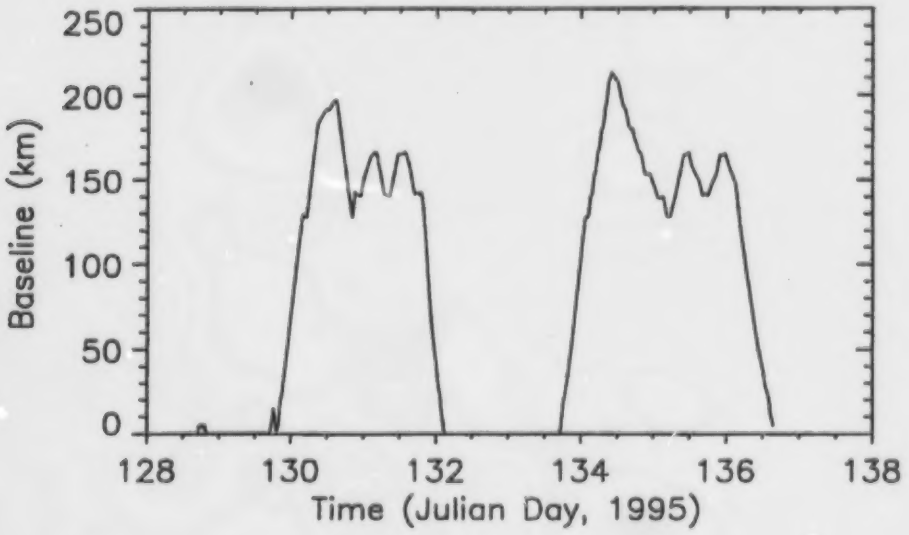

Everything seems to point to dGPS accuracies on the order of 0.5 to 1.0 metres RMS in each position coordinate.

2.2.1. Addendum

Since this analysis was originally written, the final report from the U of C has been located (reference [5]). It includes additional information on the specific times during which dGPS reference data was collected and post-processed, as well as plots of HDOPs (typically averaging 1.4), VDOPs (typically averaging 2.0) and number of satellites (typically averaging between 6 and 7) used during

processing, plots of ship trajectories, and tables summarizing the least squares residuals calculated for each satellite for each processing run. An accuracy assessment was also presented. Using carrier phase measurements in the U of C's FLYKIN™ software with fixed integer ambiguities over selected segments of two sessions, RMS coordinate differences of between 0.3 and 0.8 metres were calculated. A 2 hour plot of height differences showed a maximum of about 3.5 metres with a mean of 0.06 and RMS of 0.72 metres. These new results confirm the assessment made above.

Chapter 3: NAVIGATION SENSORS: DESCRIPTION AND ANALYSIS

This chapter looks at each of the navigation sensors used in this DIINS sea trial:

1. Two Sperry MK29 Mod 3 INSs;
2. One Trimble TASMAN P(Y) code GPS receiver;
3. One Internav LC360 Loran-C receiver.

A brief general description is followed by an analysis of the sensor's performance during ST95.

3.1: Trimble TASMAN P(Y) Code GPS Receiver

The Trimble TASMAN P(Y) GPS receiver is a military system which is capable of decrypting the GPS (encrypted) Y code to optimize stand-alone GPS performance. It is the standard issue GPS set on the Canadian Navy's latest frigates (Canadian Patrol Frigates — CPFs). All TASMAN data used and collected during ST95 were generated using Y code data.

3.1.1. Data Formats

Data can be output from the TASMAN in one of two formats:

1. TIPY is a Trimble proprietary format. With its two way communication capabilities, a large number of parameters in the receiver can be set by the user (eg. initial time and position) and/or transmitted to the user (eg. ephemerides, expected position accuracy, DOPs, etc.).
2. NMEA 0183 is an interfacing standard established by the National Marine Electronics Association. The TASMAN user is able to select a limited number of output records in NMEA format through the TASMAN's NavTrak™ Control and Display Unit (CDU). Only four record types are useful in DIINS processing:
   a. "GGA" — contains GPS latitude and longitude, a status flag, the number of satellites used to compute the position and HDOP.

b. "VTG" — contains track (course over ground, COG) and speed (speed over ground, SOG).

c. "ZDA" — contains the present (UTC) time and date.

d. "GLL" — contains GPS latitude and longitude. (This record is not used by DIINS; it is distributed to other users on the CPFs.)

Notice that there is no measure of position accuracy from the NMEA records. This is the most significant deficiency in the NMEA output. NMEA 0183 Standard Version 2.1 (and possibly earlier versions) specifies a GPS quality indicator in record type GGA which identifies the source of GPS fixes as

1. Not available or invalid,
2. Valid SPS mode,
3. Valid differential, SPS mode,
4. Valid PPS mode.

This information would be very useful by enabling the DIINS filters to switch to the most appropriate error model whenever the receiver's fix mode changed. For example, if the receiver's encryption keys expired, the error model should change from PPS to SPS. But, the NavTrak™ CDU's GGA record does not follow the standard and provides only a valid / invalid status flag. Thus, there is no way to communicate the change to DIINS. Discussions with Trimble's Military Systems division indicated that NMEA records are also available directly from the TASMAN receiver through the "Maintenance Port". The GGA record from this source does follow the NMEA standard and uses the full quality indicator. Use of this option is currently being investigated (9 October 1996).

The position accuracy information available from the TIPY option is very helpful for DIINS' filters. There are two estimates from the TIPY record which provide internal receiver position accuracy estimates: EHE (Estimated Horizontal Error) and EVE (Estimated Vertical Error). Another phone call to Trimble confirmed that EHE is the RMS of the latitude and longitude standard deviations (dRMS). This information is even more useful than the NMEA quality indicator because it changes as the expected accuracy changes and does not rely on generic, a priori error models. Note: the expected position error (one sigma) of either horizontal coordinate is $EHE/\sqrt{2}$. All DIINS sea trials with the Trimble TASMAN receiver used the TIPY data for real time as well as post- processing.

Figure 7 TASMAN Latitude Errors vs dGPS with TASMAN Predicted Bounds, Days 131–132

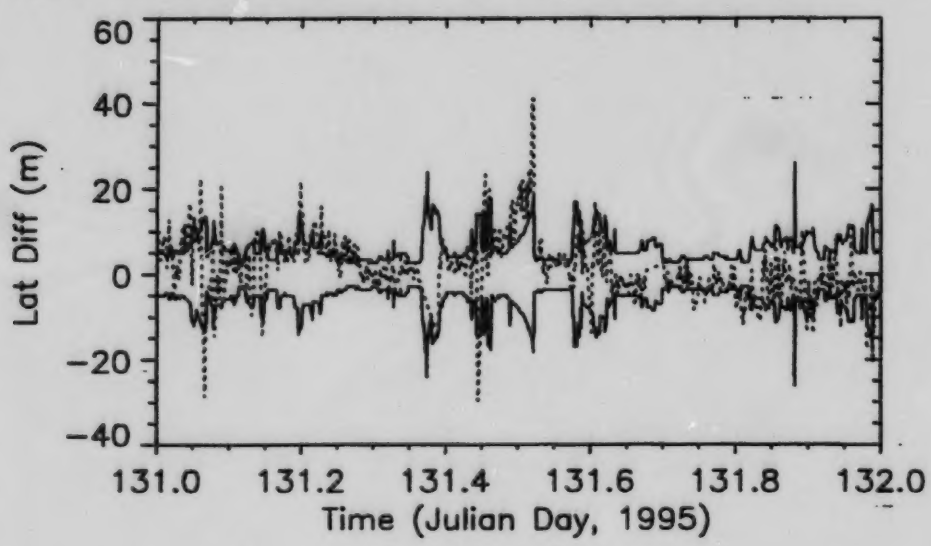

Figure 8 TASMAN Longitude Errors vs dGPS with TASMAN Predicted Bounds, Days 131–132

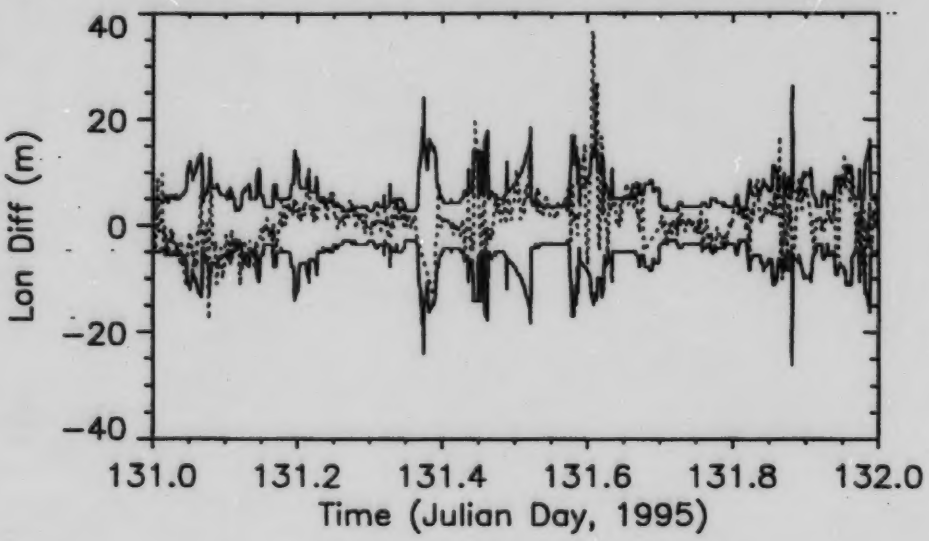

Figure 9 TASMAN Height Errors vs dGPS with TASMAN Predicted Bounds, Days 131–132

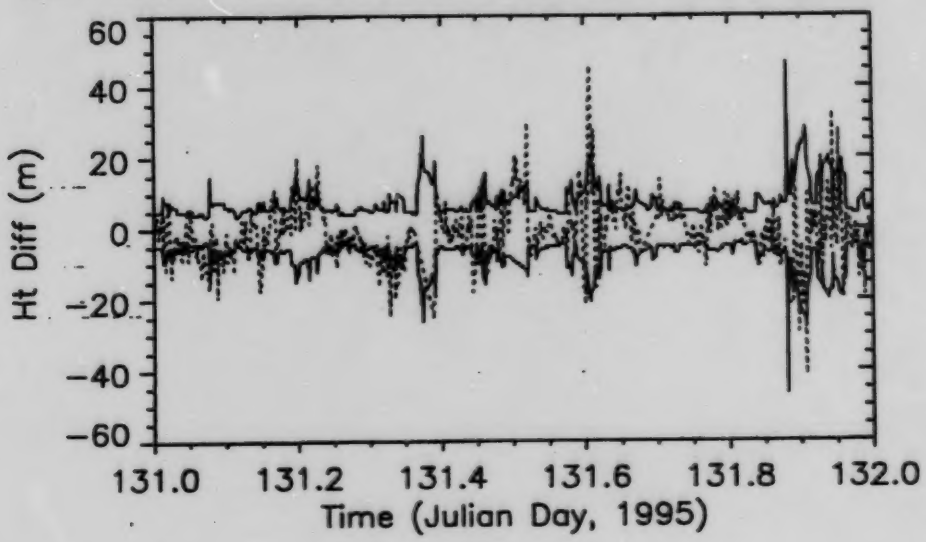

Figure 10 TASMAN Latitude Errors vs dGPS with Expected Bounds, Days 135–136

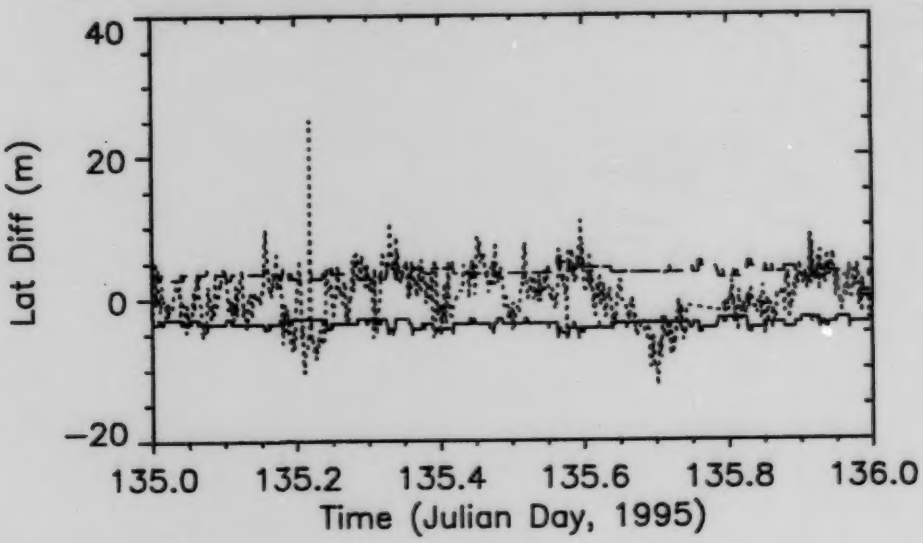

Figure 11 TASMAN Longitude Errors vs dGPS with Expected Bounds, Days 135–136

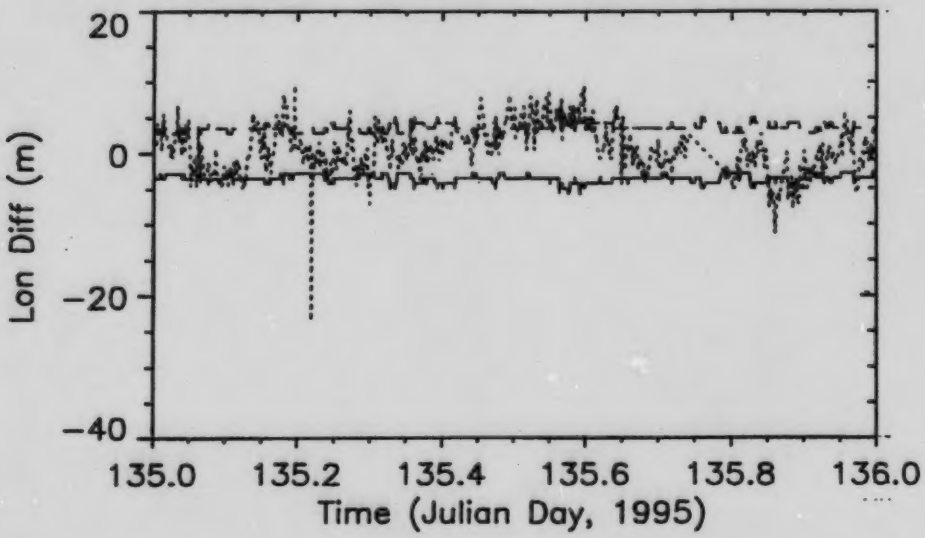

Figure 12 TASMAN Height Errors vs dGPS with TASMAN Predicted Bounds, Days 135–136

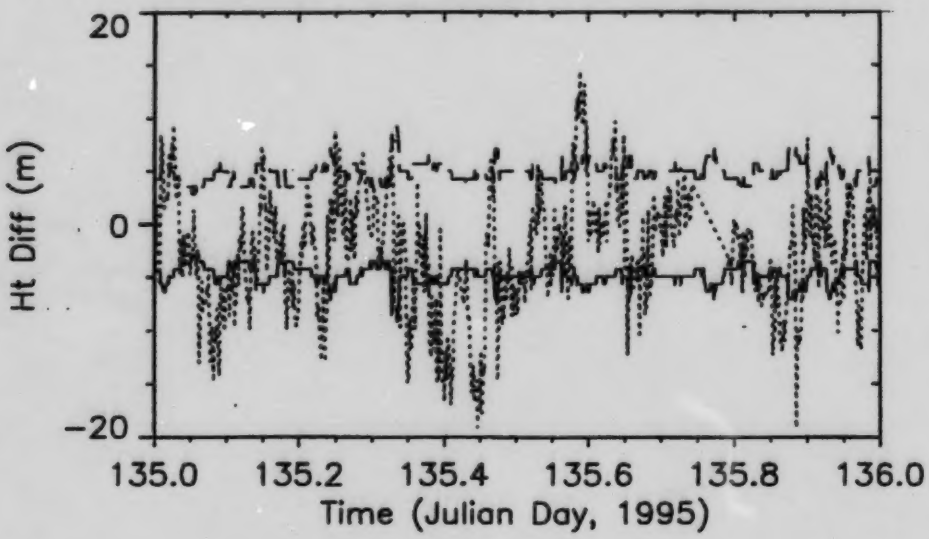

Figures 7 to 12 are plots of TASMAN position errors versus the dGPS reference as well as the 1–sigma bounds expected by the TASMAN receiver. Horizontal bounds are $\pm EHE/\sqrt{2}$. Vertical bounds are $\pm$ EVE. These days were chosen to represent the full run. The TASMAN data has higher levels of noise over days 129–133 and lower levels over days 134–136. The reason for this difference is not clear: it may be atmospheric, different levels of GPS selective availability, etc.

3.1.2. Position Errors vs Differential GPS

All four data sets (Days 129–132, 133, 134–135 and 136) were combined to generate the following position coordinate statistics:

Table 5 dGPS — P(Y) Position Differences

Coordinate	Average (m)	RMS (m)	Std.Dev. (m)	Maximum (m)
Latitude	0.63	5.66	5.62	50.6
Longitude	0.42	4.70	4.68	58.1
Height	-1.08	7.79	7.72	-92.6

Assuming dGPS contributes about 1 metres RMS to each coordinate, the RMS error in the TASMAN P(Y) positions is about (5.6, 4.6, 7.7) metres, respectively. These results were generated using comparisons of 411,041 out of 463,695 (52,654 or 9% were rejected). In this run, all data which met the maximum time difference criterion (in this case, it was 0.5 seconds) was interpolated to the P(Y) time. Recall that the time difference criterion used in the data gap analysis was 10 seconds.

Figures 13 to 15 show the dGPS / TASMAN P(Y) position differences. The large gap in the centre of the plots is the time between Run 129–132 and Run 133 when the ship was in port at Esquimalt. The "outliers" are not visible on the plots because only every 50th point was plotted.

Figure 13 TASMAN P(Y) Latitude Differences vs dGPS

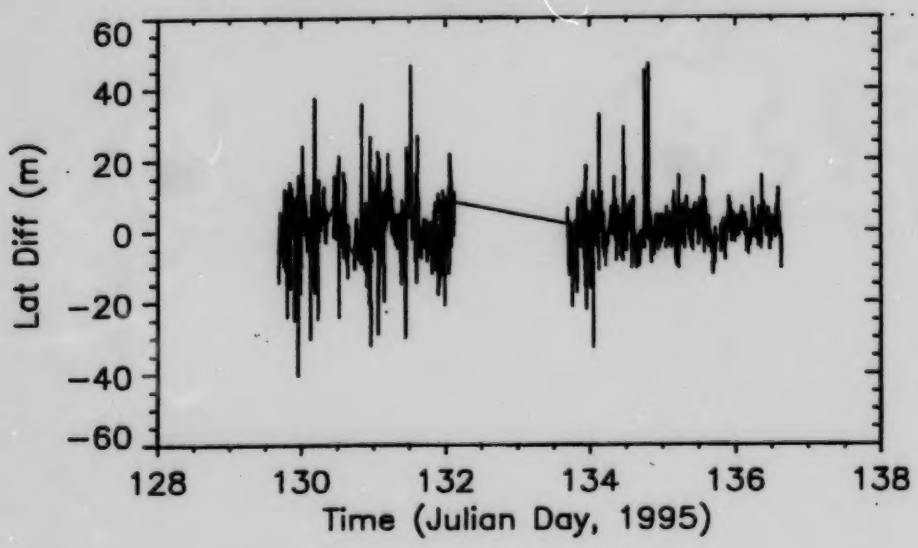

Figure 14 TASMAN P(Y) Longitude Differences vs dGPS

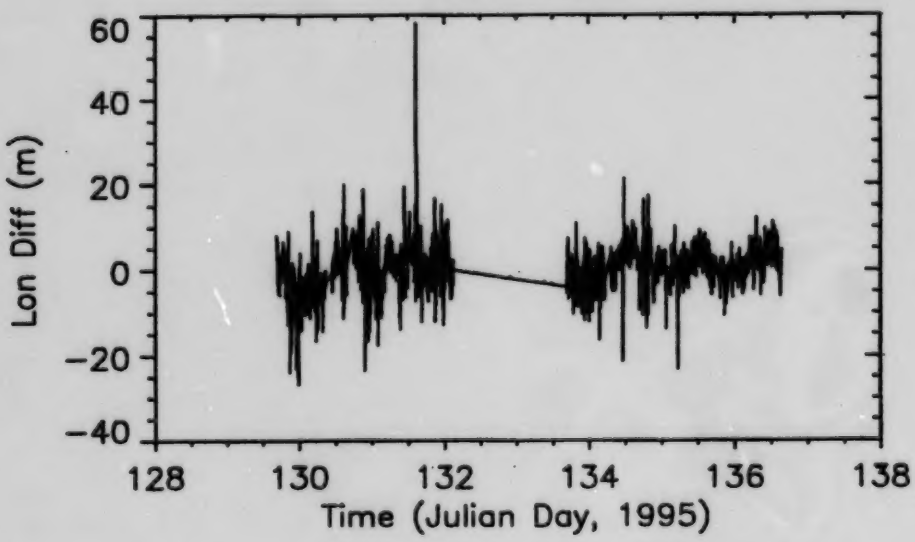

Figure 15 TASMAN P(Y) Height Differences vs dGPS

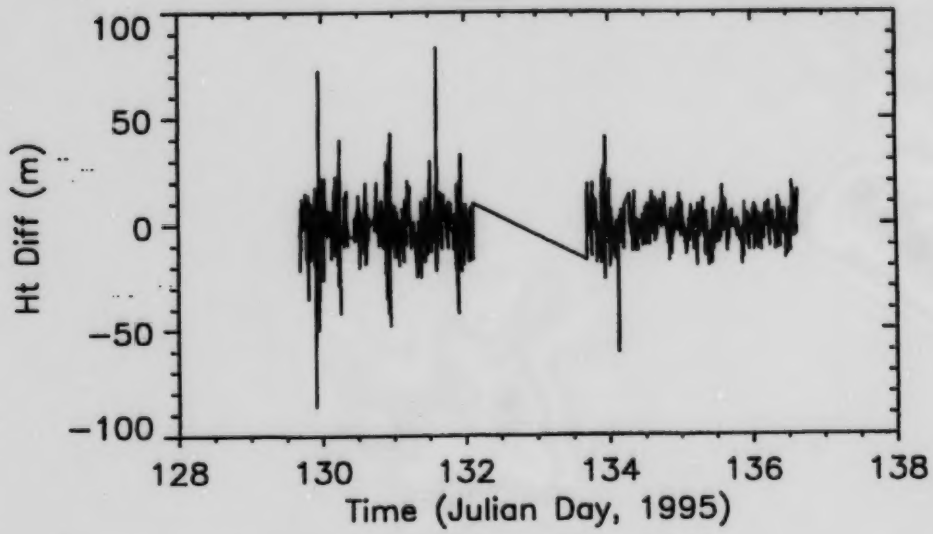

By removing all position differences which exceed a 5–sigma bound (where running statistics are computed after every 100 data points) and a 30 metre hard upper bound, 4,493 outliers were detected. The summary under these conditions is presented in the following table:

Table 6 dGPS — P(Y) Position Differences

Coordinate	Average (m)	RMS (m)	Std.Dev. (m)	Maximum (m)
Latitude	0.63	5.13	5.09	30.0
Longitude	0.44	4.45	4.43	-29.1
Height	-1.07	6.79	6.71	-30.0

406,548 data points out of 463,695 were compared; 57,147 (12%) were rejected as outliers.

3.1.3. Filter Models

All DIINS filters using TASMAN GPS processed position measurements and included 2 GPS position states and a GPS time lag state. The GPS measurements

are the most accurate available using this configuration of aiding sensors: the position states are used primarily to establish a reasonable minimum level for the filter estimated position standard deviations. The time lag state is used to estimate any time bias which may present between the INS data and GPS data. Models were developed using stationary data collected at DREO, earlier sea trial data with differential GPS as position reference and data collected during van tests in the vicinity of DREO.

The results of the stationary data analysis are summarized in Table 7. 6,861 data points, one every 10 seconds, were collected. The "true" position is assumed to be the mean of all data points. The first column contains the standard deviation around that mean. After removing the mean from each coordinate, an autocorrelation analysis was undertaken. Using simple graphical analysis methods, the standard deviation (SD) and correlation time (CT) of an assumed a first-order Gauss-Markov process along with the noise were estimated.

Table 7 Stationary P(Y) Position Errors

Coordinate	Std.Dev. (m)	Markov SD (m)	Markov CT (s)	Noise (m)
Latitude	7.6	7.4	150	1.8
Longitude	6.3	6.1	150	1.3

Next, results from the 1994 DIINS sea trial (ST94) were analyzed. In this case, the TASMAN P(Y) data was compared with the differential GPS reference. 160,439 data points, one every second, from about Day 322.7 to 325.5 were computed and used to provide the results summarized in Table 8. The mean error

Table 8 P(Y) Position Errors vs dGPS, ST94

Coordinate	Average (m)	RMS (m)	Std.Dev. (m)
Latitude	1.01	4.31	4.19
Longitude	-0.48	3.81	3.78

was subtracted from each coordinate to allow an autocorrelation analysis. The

results are summarized in Table 9. Notice that the standard deviations are smaller

Table 9 P(Y) Position Errors vs dGPS, ST94 — Gauss-Markov Model

Coordinate	Markov SD (m)	Markov CT (s)	Noise (m)
Latitude	4.0	1100	1.2
Longitude	3.6	3800	1.0

than for the stationary tests. If any difference was expected, the less benign environment for the stationary test (no motion, no lever arms, etc.) should have produced smaller errors not larger. It is possible that local multipath, satellite shading or other effects may have affected the stationary test — any position biases due to such effects would be absorbed by the standard deviations (recall that the mean errors were zero by default in the stationary analysis). Note also that the correlation times are much longer for the sea trial data. This suggests that the P(Y) and dGPS position errors may be correlated.

The same analysis was undertaken using 1995 sea trial data. The error summaries were given in Table 5. As usual, the mean error was subtracted from each coordinate to allow an autocorrelation analysis. The results are summarized in Table 10. Again, the standard deviations are smaller and the correlation times

Table 10 P(Y) Position Errors vs dGPS, ST95 — Gauss-Markov Model

Coordinate	Markov SD (m)	Markov CT (s)	Noise (m)
Latitude	5.3	800	1.8
Longitude	4.5	> 5000	1.4

are longer than for the stationary tests, reinforcing the conclusions drawn from the 1994 data set.

Finally, the TASMAN P(Y) GPS receiver was mounted in mobile lab on a flat bed truck for additional testing. An Ashtech GPS receiver was used as the

differential GPS reference. The mean error was subtracted from each coordinate

Table 11 P(Y) Position Errors vs dGPS, Van Tests August 1994

Coordinate	Day	Average (m)	RMS (m)	Std.Dev. (m)
Latitude	234	1.20	5.58	5.45
	242	-1.98	6.98	6.70
Longitude	234	-2.75	5.78	5.09
	242	-2.65	4.95	4.18

to allow an autocorrelation analysis. The results are summarized in Table 12.

Table 12 P(Y) Position Errors vs dGPS, Van Tests 1994 — Gauss-Markov Model

Coordinate	Day	Markov SD (m)	Markov CT *(s)	Noise (m)
Latitude	234	4.9	144-240	2.3
	242	6.0	126-210	2.8
Longitude	234	4.6	174-290	2.2
	242	3.8	210-350	1.7

The Markov model from the van tests is not very reliable because the data points are not very evenly spaced. They range from 1 second to 70 seconds with an average of 3–5 seconds. This is the range used to compute the correlation times in Table 12.

How do the models compare between the four data sets described above? Consider the following comparisons.

1. Latitude vs longitude:
    a. Latitude errors are almost always larger than longitude;
    b. Latitude correlation times are shorter than longitude —
        - except for the stationary data where they are equal,
        - especially for sea trial data.

2. ST94 vs ST95:

   a. ST94 errors are smaller than ST95;
   b. Correlation times are similar, but —

      - ST94 latitude correlation time is longer than ST95;
      - ST94 longitude correlation time is shorter than ST95.

3. Stationary vs van tests:

   a. Stationary Gauss-Markov standard deviations are larger than van tests;
   b. Stationary noise is smaller than van tests;
   c. Correlation times are similar, but —

      - Latitude and longitude are the same for the stationary data;
      - Longitude has longer correlation times for van tests.

4. Land vs Sea:

   a. Correlation times are shorter on land (stationary tests and van tests) vs sea (ST94 and ST95).

5. Estimated models vs expected models:

   a. Error magnitudes agree fairly well with expected models.
   b. Correlation times for the stationary data are within expected limits. The van tests' correlation times are probably unreliable. The sea trials' correlation times are longer than expected.

Why would the TASMAN P(Y) latitude errors vs the dGPS reference tend to be larger than longitude errors? The differential GPS reference data supplied by the University of Calgary for ST94 includes estimates of latitude and longitude standard deviations. These indicate that latitude accuracy is expected to be lower than longitude. However, N(orth)DOPs provided by the U of C for ST95 are generally smaller than the E(ast)DOPs, indicating that satellite geometry is generally stronger in the north-south (latitude) direction. All else being equal, with these DOPs, latitude would have higher expected accuracy. However, if satellites mainly north or south of the receiver had larger pseudorange errors vs satellites mainly east or west, the expected latitude accuracy could still be lower than that for longitude. Larger pseudorange errors could be due to larger atmospheric effects or increased multipath which in turn could be the result of lower satellite elevations. However, it is not clear why satellites mainly north and

south would be expected to be lower than satellites mainly east and west. Since the stationary data exhibits these differences, they are likely not due to P(Y) vs dGPS effects.

The large differences in estimated correlation times between latitude and longitude are difficult to explain. Since this is not observed in the stationary data, it is most likely due to some P(Y) vs dGPS effects. The most likely source is correlation between errors from the two systems. Of the common GPS errors, tropospheric refraction, multipath, and possibly ephemeris errors could be the culprits.

Given the similarities between ST94 and ST95 — same geographic location, same receivers, etc., it is not clear why ST94 errors would be smaller than ST95 errors. Seas were generally rougher during ST94 and the distance between the two antennas was longer. Since both of these effects would tend increase the computed position errors, lever arm effects can be confidently eliminated as a possible cause. The differences in correlation times from ST94 to ST95 cannot be explained — it may simply be modelling uncertainty. The differences from ST94 to ST95 are much smaller than differences between latitude and longitude.

The differences in the stationary data models and the van test models are probably within modelling uncertainties.

Since the correlation times estimated from both sea trial data sets are significantly different from the land based models and also those anticipated from past experience and the literature, the sea trial models are assumed to be unreliable. The model used for the Trimble TASMAN P(Y) code receiver during the DIINS 1995 sea trial is as follows:

1. Measurement noise = 2 metres,
2. First order Gauss-Markov standard deviation = 7 metres,
3. First order Gauss-Markov correlation time = 120 seconds.

Conservative estimates of accuracy were selected. Correlation times from the stationary data were judged to be the most reliable but a previously estimated value was used. Based on the above analysis, a correlation time of 150–200 seconds might be more appropriate.

3.2: Sperry MK29 Inertial Navigation Systems

The Sperry MK29 Mod 3 INSs used by the Canadian Navy are gimballed, stabilized platforms. The accelerometers are maintained in a north-south, east-west orientation. There is no vertical channel mechanized in the Canadian systems. There are two two-degree-of-freedom tuned rotor gyros. The gyro which senses rotations about the horizontal axes (ie. with its spin axis nominally vertical) is "indexed" (ie. rotated by 180° about its spin axis) every 210 seconds to average out some gyro errors and thus improve long term accuracy of the gyros. This, in turn, improves the levelling accuracy of the stable platform. The MK29s are also equipped with fixed gain velocity damping filters (one for north and one for east) based on rectified speed log data. These filters are designed to reduce Schuler oscillations in the INSs. The filters can be activated or deactivated manually by the operator or automatically via software. The Sperry Kalman Optimal Reset (SKOR) filter uses position resets from a number of possible sources in a 10 state Kalman filter to control position errors. Corrections are applied in a closed loop fashion. In all DIINS sea trials, including ST95, manual position resets were applied: the operator notes the current position from the nearest source (ie. a nearby GPS receiver) and manually enters that position along with an estimate of its accuracy via the MK29's control panel. The operator must specifically accept every manual fix — the filter resets are not applied until the fix is accepted. Fixes can also be entered automatically through the Serial Digital Interface with positions from an independent navigation sensor (eg. GPS, Loran-C, Omega). The SKOR filter then tests the fix for consistency (ie. a performs a residual test) and is capable of rejecting bad fixes. The operator may override an automatic fix rejection.

3.2.1. Position Errors

To analyze the INS errors, the Trimble TASMAN (P(Y) code GPS positions were used. The TASMAN data has been selected because

1. It matches the MK29 data better — data was collected from all DIINS sensors whenever DIINS was running.
2. The INS errors are much larger than the GPS errors — it makes little difference to the analysis if the reference is accurate to 1 metre or 5 metres.

Figure 16 MK29-1 Latitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds

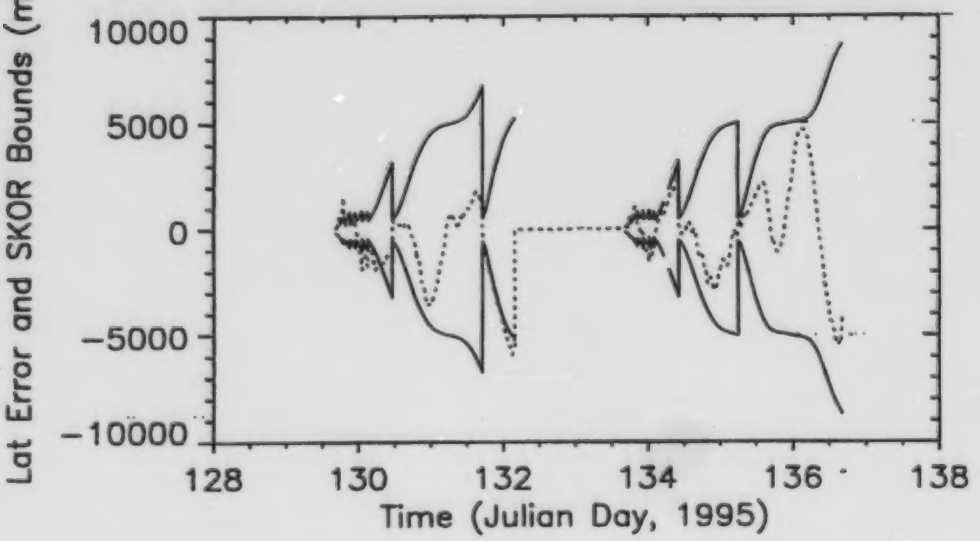

Figure 17 MK29-1 Longitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds

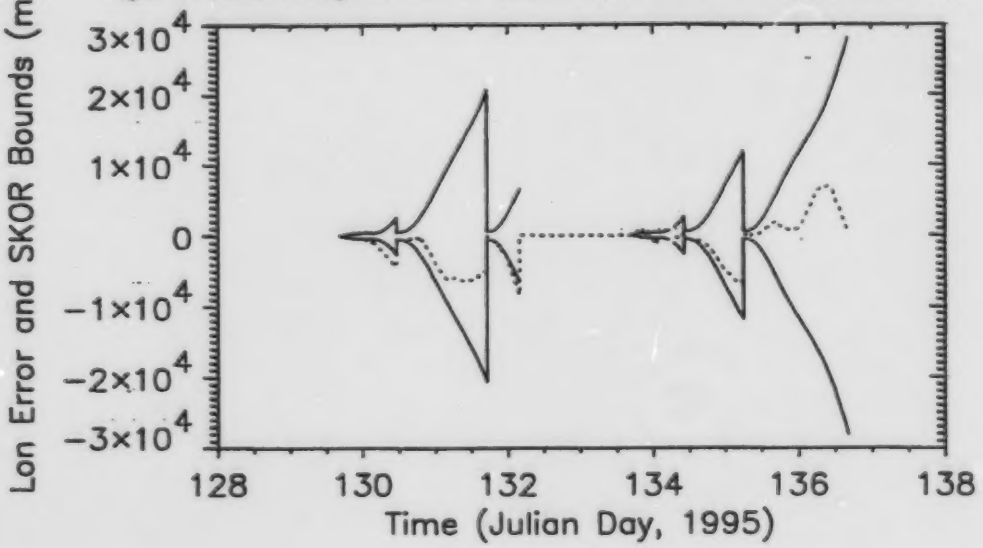

Figure 18 MK29–2 Latitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds

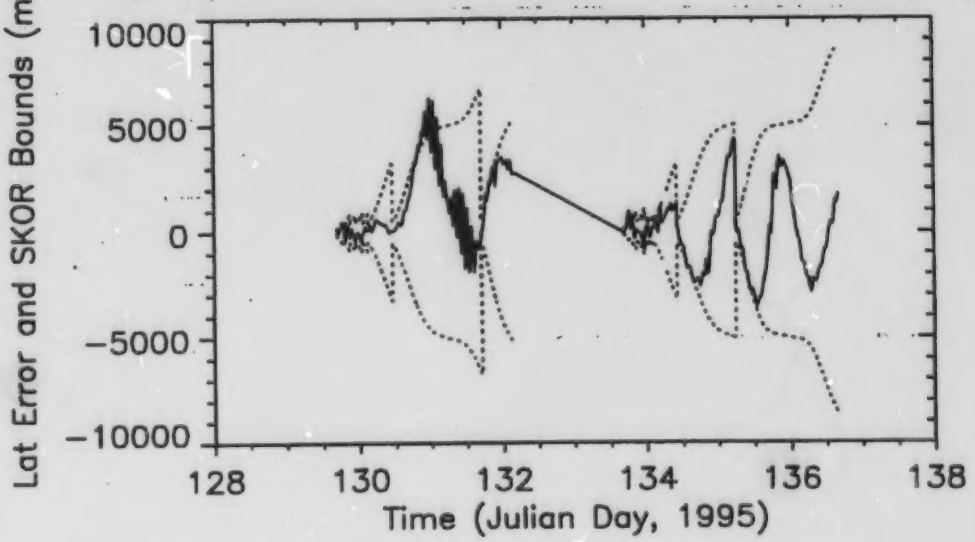

Figure 19 MK29–2 Longitude Errors vs TASMAN P(Y) with SKOR Predicted Bounds

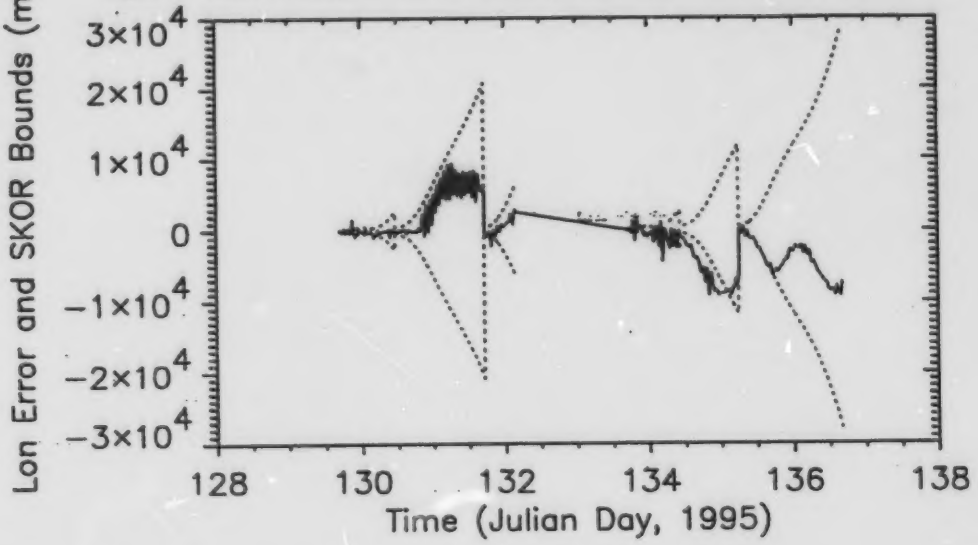

Figures 16 to 19 show errors in latitude and longitude for both INSs (the MK29s do not output height data). The Sperry MK29s have an internal Kalman filter (the Sperry Kalman Optimal Reset — SKOR filter) which processes position resets from an independent source. The latitude and longitude standard deviations predicted by the SKOR filter are plotted with the computed position errors.

Free running INS latitude and longitude errors oscillate at the *Schuler* period (about 84 minutes) and diurnally. Latitude errors are bounded, longitude errors are not. Being a marine system, the MK29s must run for days, weeks or maybe even months at a time. To prevent excessive position error growth (thus keeping the INS error behaviour within modelled bounds), the MK29s are given periodic position resets. During ST95, manual position resets were entered by the operators at various rates. The following table summarizes the reset schedule used.

Table 13 ST95 INS Position Reset Schedule

Start time (JJJ:hh:mm)	End Time (JJJ:hh:mm)	INS Reset Interval (hours)
129:16:16	130:03:10	1.3 - 2.9
130:03:10	130:11:17	8
130:11:17	131:17:04	30
Dockside		
133:16:29	134:02:12	0.8 - 2.9
134:02:12	134:10:17	8
134:10:17	135:05:54	20
135:05:54	136:16:18	34

Reset times are given for the forward INS(1), but are very similar for the aft INS(2).

The Sperry MK29 INSs also use fixed gain third order loops to filter velocities with the output from the ship's speed log. There are independent damping filters for north (south) and east (west) velocities. These "damping" loops are normally activated and deactivated automatically by the INS. However, the operator can manually override the automatic mode selection. Activation of velocity damping results in large attenuation of the Schuler errors. Figures 20 shows an example of

error behaviour with and without velocity damping. The velocity damping loop is mostly ON except for the period of time from about Day 130.8 to Day 131.6.

Figure 20 MK29 Position Errors with Velocity Damping On and Off

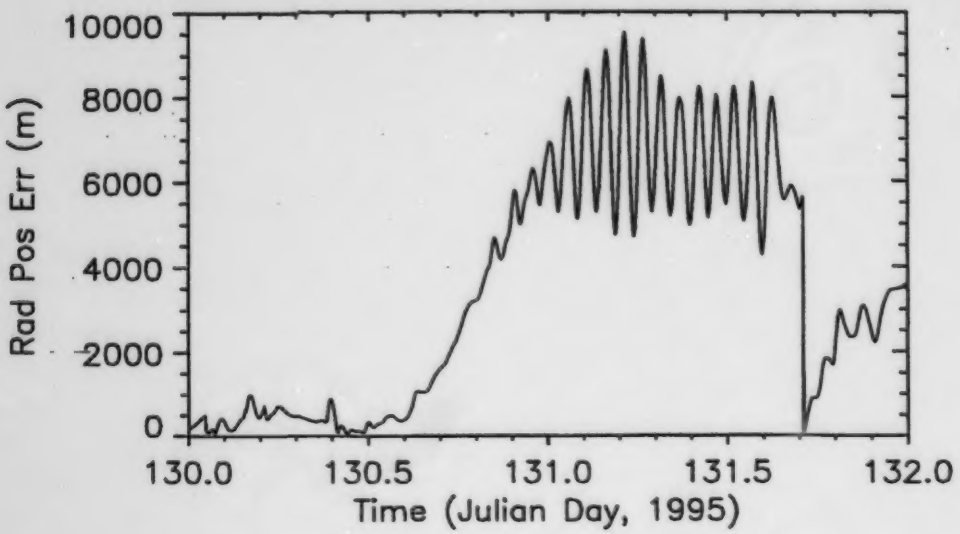

3.2.2. SKOR Manual Fix Accuracy

As discussed above, the MK29s run a SKOR (Kalman) filter to allow the resetting of selected inertial parameters: position, heading, gyro errors and perhaps velocity. The SKOR filter accepts position fixes manually or through the SDI. All DIINS sea trials have used manual, operator entered fixes. For ST95, the following specific steps were used to implement each manual fix:

1. The ship's position was read from the display of a nearby Trimble P(Y) GPS receiver — there was one receiver mounted near each INS.
2. The GPS position was manually entered into the MK29 via the control panel.
3. The fix's estimated standard deviation was entered via the control panel. A value of 0.3 nautical miles (about 550 metres) was used for every fix. It can be inferred from the data that this is the level of measurement noise used for *each* coordinate (see the discussion below).
4. Within a few seconds, the MK29 processed the fix and displayed some of the updated SKOR estimates of its states and standard deviations. If the new

latitude and longitude standard deviations were both less than or equal to the input fix standard deviation (0.3 nautical miles), the fix was accepted.

To judge the efficacy of the manual fixes accepted during ST95, Table 14 was compiled. It shows the approximate times of all ST95 fixes, the radial position error before the fix (for each INS) and the radial position error after the fix (for each INS). Note the large errors found immediately after some of the fixes were applied. Assuming the fix errors were dominated by delays in transferring the GPS position to the MK29 control panel (at typical speeds of 10 knots or 5 metres per second) and quantization errors (the manual fix accuracy is limited to 0.01 arc minutes — about 19 metres and 12 metres for latitude and longitude, respectively), manual fixes were expected to be accurate to within 100 metres. A more detailed analysis below shows that this does reflect the accuracy of a typical fix.

It is not clear from any available Sperry references whether the "sigma" input for the fixes (0.3 nautical miles for all ST95 fixes) is the standard deviation of each coordinate or an RMS value. However, this can be inferred by looking at the position standard deviations from a SKOR filter immediately before and after a particular kind of fix. The SKOR filter processes a fix by combining its pre-fix estimate of the position error and the measurement (which is the difference between the reference position and the INS computed position). The pre-fix SKOR position error and the measurement is each assigned a weight based on the inverse of its variance (a simplified Kalman filter is described below). Thus, if the SKOR variance before the fix is much larger than the measurement variance, the post-fix SKOR solution will be heavily weighted towards the fix information and lightly weighted towards the pre-fix SKOR information. As a result, the new INS position will be driven close to the reference position and the post-fix SKOR variance will be close to the fix variance. Prior to 135.24 days, the SKOR filters had run for 0.8 days without a fix. The MK29–2 SKOR latitude standard deviation grew to about 5000 metres; longitude grew to 12,000 metres. At 135.24, a fix was entered and accepted with a standard deviation of 0.3 nautical miles (or about 550 metres). This is much smaller than both pre-fix SKOR position standard deviations. After the fix, the SKOR latitude and longitude standard deviations had both dropped to 550 metres (very close to the "sigma" input for the fix). This leads to the conclusion that the fix "sigma" is used as the standard deviation of each coordinate. If the input fix "sigma" was an RMS value, each coordinate would have a standard deviation of about 390 metres (550 metres divided by

Table 14 Radial Position Errors Before and After All ST95 Fixes

Time (Run#)	MK29-1 Before Fix (m)	MK29-1 After Fix (m)	MK29-2 Before Fix (m)	MK29-2 After Fix (m)
129.79 (1)	1585	760	200	150
129.88 (1)	950	280	1460	1040
129.96 (1)	400	280	510	130
130.05 (1)	2070	490	500	120
130.13 (1)	1990	480	220	160
130.47 (1)	4480	720	110	110
131.70 (1)	5220	50	5620	50
132.16 (1)	10,020	20	—	—
133.77 (2)	560	350	1050	680
133.86 (2)	1100	480	80	30
133.94 / 133.99 (2)	590	150	1190	520
134.02 (2)	1800	520	770	550
134.09 (2)	1470	550	370	230
134.43 (3)	2640	300	1525	370
135.24 (3)	6970	100	8220	220

$\sqrt{2}$). MK29–1 SKOR results confirm this conclusion. All relevant data can be found in Appendix A, at 74,757,073.0 seconds and 74,757,225.7 DIINS time for MK29–2 and MK29–1, respectively.

Let's try to find out why some of these post-fix errors are so large. Figures 21 and 22 show the MK29–2 latitude and longitude errors, along with the SKOR standard deviations, across the fix at 129.88 days. This is the fix which had a radial position error of 1040 metres after the fix. Take note of the following:

1. According to Table 15, the error in the entered fix was 20 metres in latitude and 140 metres in longitude. Both errors are well within the 0.3 nautical mile (550 metre) fix standard deviation entered by the operator and way below

Figure 21 MK29-2 Latitude Errors Across a Reset with SKOR Predicted Bounds

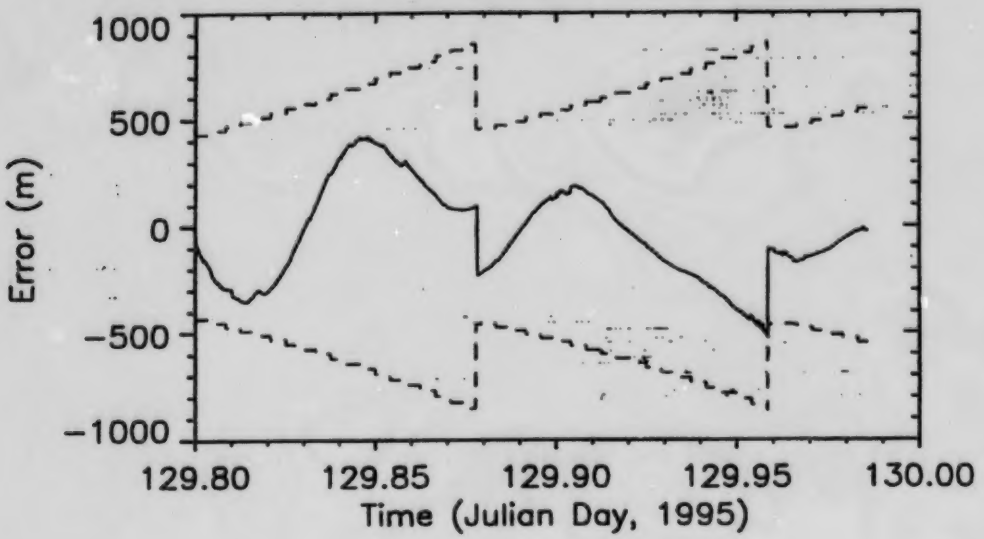

Figure 22 MK29-2 Longitude Errors Across a Reset with SKOR Predicted Bounds

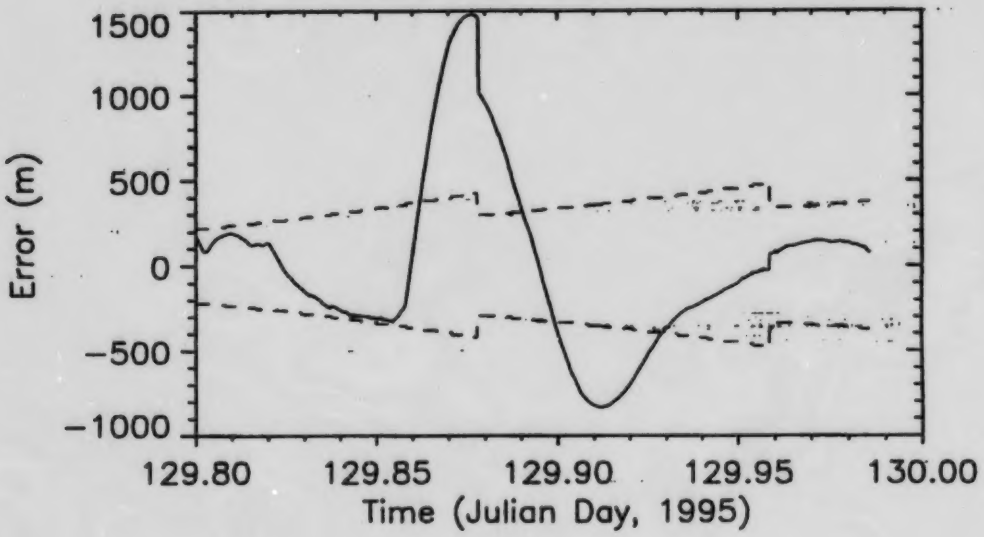

the post-fix errors of 280 and 1040 metres, respectively. Therefore, we can conclude that it was not a bad fix which caused the large post-fix errors.

2. Longitude errors immediately prior to the fix are much larger than latitude errors (1460 metres vs 90 metres) — the radial error is dominated by the longitude error.
3. The SKOR longitude standard deviation immediately prior to the fix is smaller than latitude (420 metres vs 850 metres).
4. Longitude errors immediately after the fix are much larger than latitude errors (1010 metres vs 230 metres). Again, the post-fix radial position error is dominated by longitude error.
5. The SKOR longitude standard deviation immediately after the fix is smaller than latitude (300 metres vs 460 metres).

The longitude error before the fix (1460 metres) is not very well represented by the SKOR standard deviation (420 metres). Thus, the SKOR filter assumed that longitude was much more accurate than it really was. When the longitude measurement was processed by the SKOR filter, the SKOR estimate of longitude was weighted too heavily vs the fix longitude. As a result, the SKOR correction to longitude was too small, and the post-fix longitude error was larger than expected. Using the information available to it, the filter also reduced its estimate of the longitude variance: the SKOR estimate is still much too optimistic after the fix — a longitude error of 1040 metres vs a SKOR standard deviation estimate of 300 metres. Notice that the magnitude of the latitude error after the fix is actually larger than before the fix. In this case, the pre-fix SKOR latitude variance was too large: the filter assigned too little weight to the latitude error state and it absorbed most of errors in the fix latitude.

Conversely, what happened at the MK29–1 fix at time 132.16 days? Here, there was a pre-fix radial position error of 10,020 metres and a post-fix error close to 20 metres:

1. According to Table 15, the error in this entered fix was 10 metres in latitude and negligible in longitude. These errors are within the limits of the display — ie. 0′.01 which is equivalent to 19 metres in latitude and 12 metres in longitude. Both errors are well within the entered 0.3 nautical mile fix standard deviation. Note that this fix was entered while the ship was tied up at the fuel jetty in Esquimalt and therefore had zero velocity.

2. Latitude and longitude errors immediately prior to the fix are both large (5460 metres and 8410 metres, respectively).
3. The SKOR latitude and longitude standard deviations immediately prior to the fix are both large (5250 metres vs 6550 metres, respectively).
4. Latitude and longitude errors immediately after the fix are both very small — about 10 metres.
5. The SKOR latitude and longitude standard deviations immediately after the fix are the same (260 metres each).

In this case, the SKOR latitude standard deviation accurately reflected the true latitude error, and the longitude standard deviation was a bit optimistic. More importantly, the SKOR position standard deviations were much larger than the fix standard deviation. In this situation, the post-fix SKOR coordinates and standard deviations would be driven very near the fix values. It seems likely that a standard deviation of about 260 metres (0.14 instead of 0.30 nautical miles) was used for this fix.

Can any improved operational procedures be developed using this information? Let's first take a very simplified look at the Kalman filter function while processing a fix. In scalar error state formulation, each fix coordinate would be processed using something like

$$\begin{aligned} \hat{x}(+) &= \hat{x}(-) + k(z - \hat{x}(-)) \\ \sigma^2(+) &= (1-k)\sigma^2(-) \end{aligned} \tag{7}$$

where $\hat{x}$ is the SKOR estimate of the error in the coordinate. (-) and (+) identify SKOR parameters immediately before and after the fix, respectively. k is the filter gain. The measurement is computed by differencing the fix coordinate and the pre-fix INS coordinate:

$$z = \hat{p}_{fix} - \hat{p}(-) \tag{8}$$

The SKOR filter runs closed loop: at each fix,

1. each updated state, $\hat{x}(+)$, is computed,
2. after the fix is accepted, the INS is reset using the updated states,
3. all reset states are zeroed.

In a (recursive) Kalman filter, zero states are propagated as zero. Thus, for the SKOR filter, $\hat{x}(-) = 0$ and

$$\hat{x}(+) = kz \tag{9}$$

The filter gain is defined as

$$k = \frac{\sigma^2(-)}{\sigma^2(-) + \sigma^2_{fix}} \tag{10}$$

Now,

$$\begin{aligned} \hat{x}(+) &= \frac{\sigma^2(-)}{\sigma^2(-) + \sigma^2_{fix}} z \\ \sigma^2(+) &= \frac{\sigma^2_{fix}}{\sigma^2(-) + \sigma^2_{fix}} \sigma^2(-) \end{aligned} \tag{11}$$

If $\sigma^2_{fix} \ll \sigma^2(-)$, then

$$\begin{aligned} \hat{x}(+) &\simeq z = \hat{p}_{fix} - \hat{p}(-) \\ \hat{p}(+) &= \hat{p}(-) + \hat{x}(+) \simeq \hat{p}_{fix} \\ \sigma^2(+) &\simeq \sigma^2_{fix} \end{aligned} \tag{12}$$

The INS position is driven to the fix position. This is most likely to occur when $\sigma^2(-)$ is very large (eg. because of a long span of time since the last fix). This was the situation for the MK29–1 fix at 132.16 days which was discussed above.

Alternatively, if $\sigma^2(-) \ll \sigma^2_{fix}$, then

$$\begin{aligned} \hat{x}(+) &\simeq \hat{x}(-) = 0 \\ \hat{p}(+) &\simeq \hat{p}(-) \\ \sigma^2(+) &\simeq \sigma^2(-) \end{aligned} \tag{13}$$

This is like no fix at all and is not likely to occur in practice. A very inaccurate position reference (eg. a sextant or Omega) would have a large σ_{fix}, but what would be the point of entering a fix which was no more accurate than the INS.

Now let's take a close look at the residual. In our simplified model, it is defined as

$$r = (z - \hat{x}(-)) \tag{14}$$

For the SKOR filter, since $\hat{x}(-) = 0$,

$$r = z = \hat{p}_{fix} - \hat{p}(-) \tag{15}$$

The residual variance is defined as

$$\sigma_r^2 = \sigma^2(-) + \sigma_{fix}^2 \tag{16}$$

The Kalman filter residual test is often implemented as

$$\begin{aligned} &\text{Pass: } |r| \leq n\sigma_r \\ &\text{Fail: } |r| > n\sigma_r \end{aligned} \tag{17}$$

where n is a specified integer.

The Sperry manual (reference [6], pages 2–38 to 2–32) is rather vague in its discussions of fix acceptance. If the fix was entered manually, the operator is advised to "*evaluate the calculated effect of accepting the fix by examining the Delta and Sigma latitude, longitude and heading values*" before accepting a fix. The same procedure is to be followed before overriding SKOR's rejection of a fix entered automatically from an aiding sensor via the SDI. No details of the evaluation procedure are given. Perhaps the above developments can be used to assist the operator in evaluating fixes.

Notice that the latitude and longitude residuals and their variances can always be computed for a manual fix (and can usually be computed for an automatic fix) using available information:

1. $\hat{p}_{fix}$ and σ_{fix}^2 are entered by the operator for a manual fix or from the reference sensor via the SDI for an automatic fix.
2. $\hat{p}(-)$ and $\sigma^2(-)$ are both available from the MK29 display panel: "ΔLAT" or "ΔLON" and "σLAT" or "σLON", respectively.

Also, notice that all required data is available before a fix is entered. Thus, it can be used to alert the operator of a potential problem before a fix is entered, allowing him to take appropriate steps to correct the problem.

Table 15 lists all fix position errors vs the Trimble P(Y) GPS set. Fix positions were taken from the MK29–1 and MK29–2 log books. Fix errors will be a combination of GPS reference errors, errors resulting from a nonzero ship's velocity and nonzero transfer delay, and errors in transcribing the fix into the

log books. Most fix errors are in the 100–200 metre range, with a couple of exceptions:

1. The 1020 metre MK29–2 latitude fix error at 129.79. From Table 14, the MK29–2 radial position error before and after the fix were 200 and 150 metres, respectively. Therefore, it seems very unlikely that a fix with a 1020 metre error was actually used — the fix latitude was probably transcribed into the log book incorrectly.
2. The 710 metre MK29–1 longitude fix error at 133.77. At the time of this fix, the ship was travelling at about 1 metres per second north and 5 metres per second west. A 140 second delay in entering the reference coordinates could explain both latitude and longitude fix errors. However, once again, the post-fix errors in Table 14 suggest that a transcription error may be the cause of this large (apparent) longitude fix error.

The test statistic,

$$\hat{m} = \frac{|r|}{\hat{\sigma}_r} \tag{18}$$

can also be computed post-mission from the stored ST95 data. A summary of all post-mission statistics and accompanying data is included in Appendix A. Values of $\hat{m}$ range from 0.0 to 2.1, distributed as indicated in Table 16. Overall,

Table 16 Post-Mission Residual Test Statistics

$\hat{m}$ Range	MK29-1 (out of 30)	MK29-2 (out of 28)
0.0 to 0.5	11	21
0.6 to 1.0	8	4
1.1 to 1.5	7	2
1.6 to 2.0	4	0
over 2.0	0	1

the MK29–2 is better behaved in this regard than MK29–1, ie. $\hat{m}$ tends to be smaller. This suggests that

1. The residual tends to be smaller — the MK29–2 error model is more representative of its true performance.
2. Or; $\hat{\sigma}_r$ tends to be larger — since $\hat{\sigma}_{fix}$ is the same for both INSs, $\sigma_{INS2}(-)$ must tend to be more pessimistic.

Table 15 Errors in Fixes Entered for ST95

Time (Run#)	MK29-1 Fix Lat Error (m)	MK29-1 Fix Lon Error (m)	MK29-2 Fix Lat Error (m)	MK29-2 Fix Lon Error (m)
129.79 (1)	100	10	1020	200
129.88 (1)	20	120	20	140
129.96 (1)	60	90	120	110
130.05 (1)	60	90	70	110
130.13 (1)	90	70	80	100
130.47 (1)	140	10	180	0
131.70 (1)	30	20	20	60
132.16 (1)	10	0	—	—
133.77 (2)	160	710	40	120
133.86 (2)	40	90	60	120
133.94 / 133.99 (2)	230	180	120	140
134.02 (2)	90	70	40	130
134.09 (2)	160	90	190	10
134.43 (3)	50	0	110	20
135.24 (3)	40	40	50	70

In considering calculated values of $\widehat{m}$ from ST95, keep in mind the following points: The standard deviation used for ST95 is about 5 times the expected value, $\widehat{\sigma}_{fix} = 550$ m $\simeq 5\sigma_{fix}$ (Table 15 indicates that most manual fixes are accurate to about 100 metres). A conservative value was selected to reduce the chance of corrupting the SKOR filter with a bad fix. The affect of this choice on the residual test depends on the relative size of $\sigma^2(-)$. Consider (from equation 16),

$$\sigma_r^2 = \sigma^2(-) + \sigma_{fix}^2 \simeq \sigma^2(-) + \frac{\widehat{\sigma}_{fix}^2}{25} \qquad (19)$$
$$\widehat{\sigma}_r^2 = \sigma^2(-) + \widehat{\sigma}_{fix}^2$$

Then,

1. If fixes are infrequent, we can expect that $\sigma^2(-) \gg \widehat{\sigma}^2_{fix}$, so that

$$\frac{\widehat{\sigma}^2_r}{\sigma^2_r} \simeq \frac{\sigma^2(-)}{\sigma^2(-)} = 1 \tag{20}$$

If

$$m = \frac{|r|}{\sigma_r} \tag{21}$$

then

$$\widehat{m} = \frac{\sigma_r}{\widehat{\sigma}_r} m \tag{22}$$

and, under these conditions $\widehat{m} \simeq m$. There will be little affect on the residual test. This is the case for almost all of the ST95 fixes.

2. If fixes are frequent, we can expect $\sigma^2(-) \geq \widehat{\sigma}^2_{fix}$ so that

$$\frac{\widehat{\sigma}^2_r}{\sigma^2_r} \simeq \frac{2\sigma^2(-)}{\frac{26}{25}\sigma^2(-)} \simeq 2 \tag{23}$$

and $\widehat{m} \simeq m/\sqrt{2} = 0.7m$. The residual test can be written $m \leq n$ where n is a given integer. If $\widehat{m} = 0.7m \leq n$ is used instead, the chances of accepting a bad fix are increased. ST95 fixes were never given at a rate higher than once every 2 hours so this case does not apply.

3. While in Dockside mode, $\sigma(-) \simeq 100$ metres. Immediately after the INS has been removed from Dockside $\sigma^2(-) \ll \widehat{\sigma}^2_{fix}$ so that

$$\frac{\widehat{\sigma}^2_r}{\sigma^2_r} \simeq \frac{\widehat{\sigma}^2_{fix}}{\sigma^2(-) + \frac{\widehat{\sigma}^2_{fix}}{25}} \tag{24}$$

Further, if $\sigma^2(-) \ll \sigma^2_{fix}$, then

$$\frac{\widehat{\sigma}^2_r}{\sigma^2_r} \simeq 25 \tag{25}$$

and $\widehat{m} \simeq 0.2m$. Or, if $\sigma^2(-) \simeq \sigma^2_{fix}$, then

$$\frac{\widehat{\sigma}^2_r}{\sigma^2_r} \simeq 12.5 \tag{26}$$

and $\widehat{m} \simeq 0.3m$. Throughout the range of possibilities, the chances of accepting a bad fix are further increased. Again, this case was not experienced during ST95.

The rejection test used by the operators during ST95 was simply:

$$\begin{aligned} &\text{Pass: } \sigma(+) \leq \sigma_{fix} \\ &\text{Fail: } \sigma(+) > \sigma_{fix} \end{aligned} \tag{27}$$

The source of this test is not clear, but it is clear that this is not a very good test. Recall that the rigorous residual test compares the difference between the measured INS error and the filter predicted error (the residual) with the expected accuracy of the difference (the residual standard deviation). However, this test does not include any measurement data: it cannot detect a bad fix. Also, the simplified Kalman update model developed above shows that this is not really a test at all. From equation 11,

$$\sigma^2(+) = \frac{\sigma^2_{fix}}{\sigma^2(-) + \sigma^2_{fix}} \sigma^2(-) \tag{28}$$

If the ST95 fix test is to fail,

$$\frac{\sigma^2_{fix}}{\sigma^2(-) + \sigma^2_{fix}} \sigma^2(-) > \sigma^2_{fix} \tag{29}$$

or

$$\frac{\sigma^2(-)}{\sigma^2(-) + \sigma^2_{fix}} > 1 \tag{30}$$

The test will approach the failure limit as σ^2_{fix} approaches zero, but based on this model, it can never truly fail.

However, if the SKOR filter did not process (ie. internally rejected) a fix, then $\sigma^2(+) = \sigma^2(-)$ with the likely result that $\sigma^2(+) = \sigma^2(-) > \sigma^2_{fix}$. So, this test would detect an internal fix rejection. However, for a manual fix, there is no automatic fix rejection (apparently there is a blunder detection test — for example, if east longitude is entered in the western hemisphere, the operator is asked to check the fix). In early days of the MK29s at DREO, very small values of σ_{fix} were used — as small as 0.01 nautical miles (18 metres).

An improvement to the analysis of fix quality can be implemented if the operator has access to the reference sensor coordinates. He can compute the residual and its variance as described above, then substitute them into equation 18. If the resulting $\widehat{m}$ is unacceptably large, the operator would be advised try

to determine the reason. The computed residual is the difference between the INS and reference sensor coordinates. If it is much larger than the SKOR filter's standard deviation estimate, it means that

1. the reference position is wrong, resulting in the coordinate difference being larger than the true INS error, or
2. the propagation of the standard deviation was poorly modelled, resulting in an optimistic estimate.

The operator should react as follows:

1. If the reference sensor has failed, it should be repaired or replaced as soon as possible. Of course, it should not be used for fixes until this is done. If it is still operating, but below expectations, $\widehat{\sigma}_{fix}$ should be increased to reflect the degraded performance.
2. It is much more likely that $\widehat{m}$ is large because of SKOR propagation problems. Since the operator can only influence the SKOR filter through the acceptance of fixes, he can try to improve the results of the fix only by changing $\widehat{\sigma}_{fix}$. The option selected is based on seeking a balance between an improved post-fix SKOR solution and the risk of disturbing the entire INS. If the operator is confident that the fix is OK as entered:
    a. He can reduce the size of $\widehat{\sigma}_{fix}$, effectively under-weighting the inaccurate SKOR position and over-weighting the more reliable fix position. This could bring the post-fix SKOR position error more into line with the post-fix standard deviations but would also increase the risk of disturbing the INS.
    b. He can increase the size of $\widehat{\sigma}_{fix}$ to reduce the size of $\widehat{m}$. However, the improvement would be less than that which could safely be achieved: the inaccurate SKOR position would be over-weighted while the more reliable fix position would be under-weighted. Also, the imbalance between the SKOR position error and standard deviations would be further exacerbated. The risk of disturbing the INS is also reduced. Of course, the surest way to avoid disturbing the INS is to eliminate the fix all together. This option offers no real benefit.
    c. He can accept the fix as is. The resulting reset will not be perfect (since the inaccurate SKOR position will still be over-weighted), but it will provide some improvement to the SKOR solution. The risk of disturbing the INS will fall somewhere between the other two options.

After an automatic fix is sent to a MK29, the SKOR filter will perform its own residual test. If this test fails, the operator still has the option to override the test and accept the fix. The above techniques should also be useful in facilitating this decision.

Further analysis of the ST95 data in Appendix A indicates that the all the large values of $\hat{m}$ correspond to INS errors which are larger than SKOR has predicted. This means that the large $\hat{m}$ values are due to σ_r values which are too small because $\sigma(-)$ is too small. The fixes were accepted using Option 2.c in all cases.

In summary,

1. The fix "sigma" entered by the operator is the standard deviation for each horizontal coordinate (and not an RMS).
2. With competent, experienced operators, manual fixes can be used with confidence — no bad fix was accepted during ST95. Large errors observed in a few fixes appear to have been caused by errors in transcribing the fix information to the log books.
3. "Poor" fixes (defined as those which result in larger than expected post-fix errors, or errors which are significantly larger than their post-fix SKOR standard deviations) are most likely caused by poor predictions in the SKOR filter (ie. when the pre-fix SKOR standard deviation is much smaller than the true error).
4. Estimates of the SKOR residual and its standard deviation can be computed before a fix is entered. Using this information, a residual test statistic can be computed to warn the operator of potential problems. In response, the operator should check the status of the reference position sensor. If the sensor in operating in a degraded fashion, $\hat{\sigma}_{fix}$ should be increased to reflect new performance expectations. If the problem is attributed to a disagreement between the INS position error and the SKOR position standard deviation, the operator can consider a careful reduction in $\hat{\sigma}_{fix}$ (thus improving the expected post-fix SKOR results but increasing the risk of disturbing the INS), or he can accept the fix as entered (providing more a balance between post-fix SKOR results and the risk of disturbing the INS).

3.2.3. Status Messages

The MK29s also output a number of status flags. Most of them are straight forward and behave as expected. However, two of them, the *Performance*

Degraded and *Inertial Well Settled* flags, are neither. Both flags switch at times to the "failure" mode. Further, the failure criteria is not well explained in the Sperry manuals. In the description of the error messages sent to the systems' displays, there is a *Performance Degraded* error. This message is sent whenever the SKOR estimated heading standard deviation exceeds 5 arc minutes or when there has been no position reset for a period of time which exceeds an adjustable limit. In the Sperry Operation and Service Manual (reference [6]), Chapter 5 Troubleshooting, Table 5–1, System Error Number Description, page 5–35, System Error Number 107 is described. The error is named "Inertial Performance Degraded" and is described as follows:

> "After first position fix is entered, the second fix must be entered within 10 hours or . . . [flag] set. When system operating in INERTIAL mode, delta heading is monitored as reference for determining accuracy. If heading deviation exceeds 5 arc minutes, . . . [flag] set. In either case, system reverts to Gyro (Fine Align) mode."
>
> The same manual (on page 2–27) states that
>
> " . . . a timer is set by the program based on the sigma value (accuracy) of the first position fix. To remain in Inertial mode, a second fix must be accepted by the system within a ten to thirty-two hour period, depending on the value of sigma entered with the fix. If the second fix is not entered before the time out period has ended, SYSTEM FAULT 107 will be announced. "
>
> Also, on the same page,
>
> " . . . the system monitors calculated standard deviation of heading (Sigma heading) value as the reference for evaluating Inertial mode performance. If the value of Sigma Heading exceeds a prescribed limit, SYSTEM FAULT 107 is displayed. "

The first two quotes state that the INS will drop out of Inertial mode (and revert to "Gyro (Fine Align) mode" when SYSTEM FAULT 107 is signalled. This suggests further questions:

1. How does this affect MK29 performance?
2. How does this impact DIINS?
3. Does the next position fix cause the MK29 to automatically resume Inertial mode operation?

The two status flags in question for both INSs are plotted in Figures 23 to 26.

Figure 23 MK29–1 Performance Degraded Status Flag

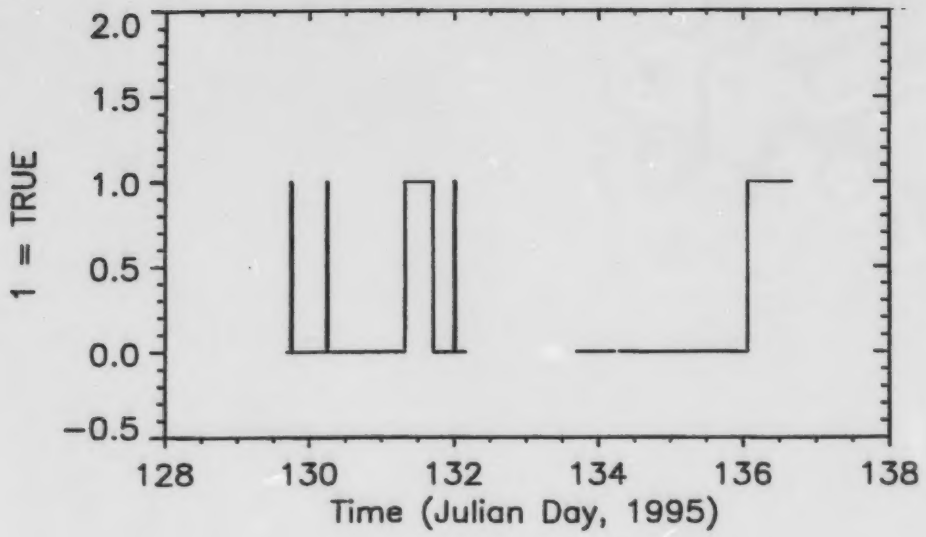

Figure 24 MK29–1 Inertial Well Settled Flag

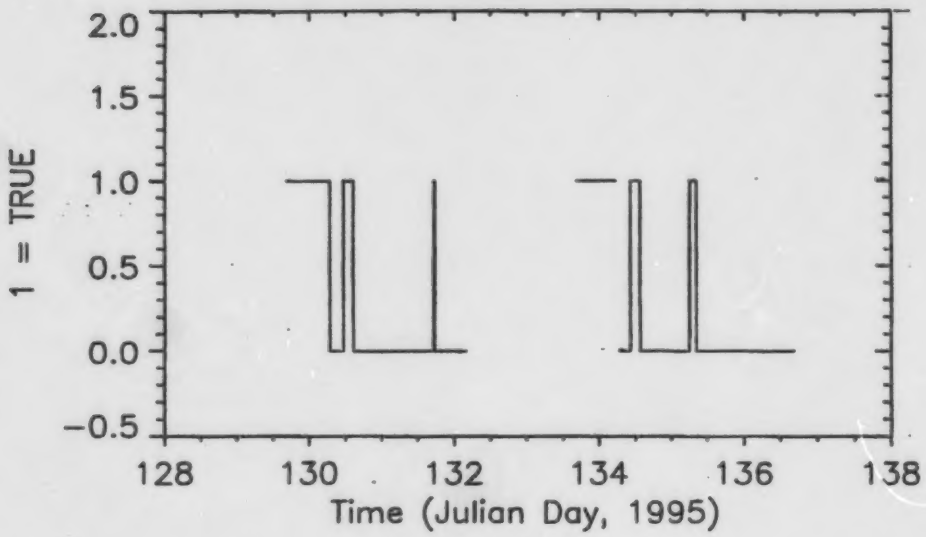

Figure 25 MK29–2 Performance Degraded Status Flag

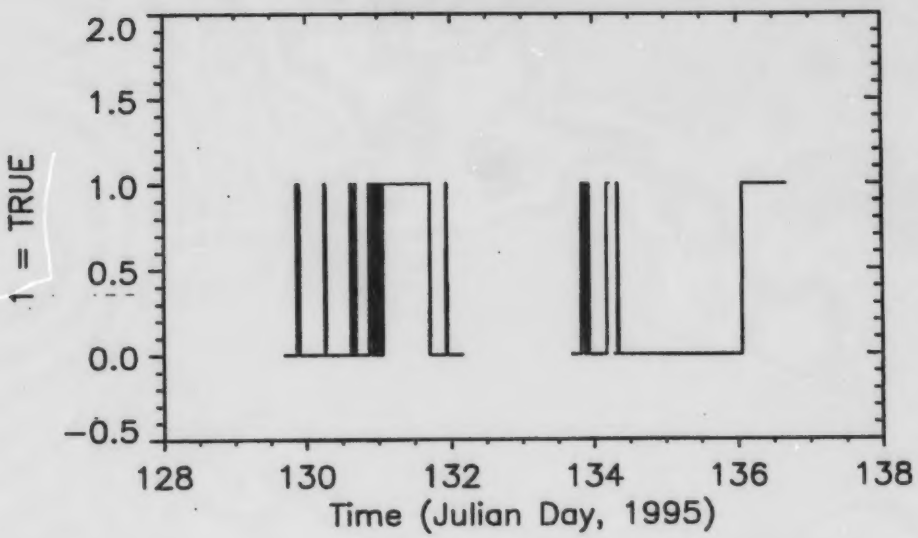

Figure 26 MK29–2 Inertial Well Settled Flag

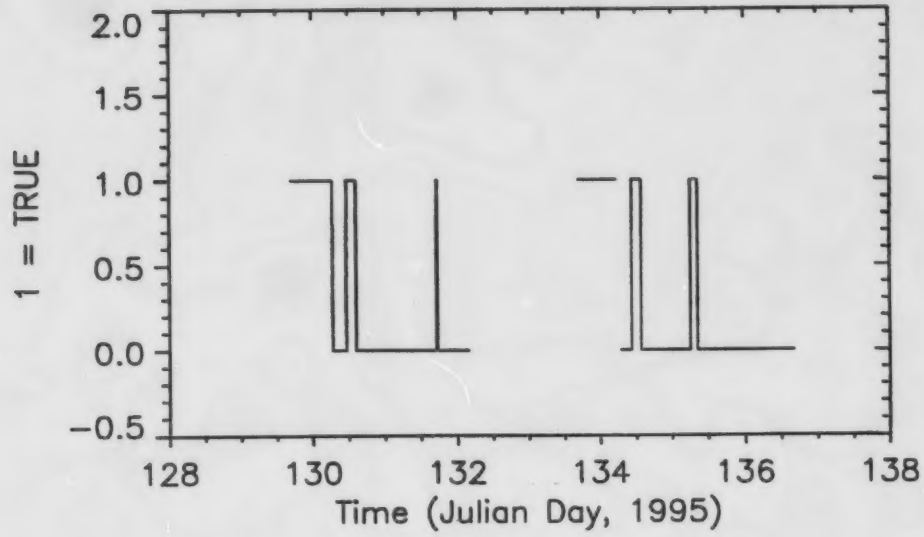

The Performance Degraded flag's being set does not seem to occur according to the criteria given. For example, it is first set at time 129:17.8. But this is less than 1.5 hours after the system was released from Dockside (ie. the last position fix). At the same time the SKOR estimate of the INS heading error standard deviation is about 1 arc minute. In this example and others, neither of the criteria quoted above for setting the flag seems to have been met. A request has been made to Sperry to supply additional information, but a reply has not yet (as of October 16 1996) been received. Nothing has been found in the Sperry manuals relating to the Inertial Well Settled flag. Empirical analysis of the data has not allowed the deduction of criteria for the setting of either of these flags: they do not seem to correspond to poorer navigation performance.

3.2.4. Filter Models

The INS error models are much too complex for analysis here. In brief, the system states (modelling errors in position, velocity and attitude) should be initialized to the best estimates of their values at start-up. Typically, all states are initialized to zero.

System state variances are initialized to the expected uncertainty at start-up. In general, it is preferable to initialize system state variances to larger rather than smaller values: the system states are driven by INS sensor error states and aiding sensor measurements. As soon as measurements are processed by the filter, the system states will be driven to values representing the accuracy of those measurements.

Sperry MK29 INS sensor errors were modelled as first order Gauss-Markov processes as follows:

1. X and Y accelerometer biases — standard deviation = 100 microg's, correlation time = 10 hours.
2. X and Y gyro biases — standard deviation = 0.001 degrees per hour, correlation time = 24 hours.
3. Z gyro bias — standard deviation = 0.003 degrees per hour, correlation time = 24 hours.

3.3: Internav LC360 Loran-C Receiver

Loran-C data was provided to DIINS by an Internav LC360 receiver. This receiver is able to output positions (as latitude and longitude) or Time Difference measurements (TDs). TDs are used in DIINS because of the added flexibility they offer in filter modelling. For example, even one TD can be used to update the filter while it could not be used to compute a position. Other advantages include improved and simpler error modelling, easier decorrelation of TD measurements versus position measurements, etc.

3.3.1. Time Difference Errors

The Loran-C TD errors off the west coast of Vancouver Island are dominated by systematic, position dependent errors. These errors are dominated by unmodelled errors resulting from delay induced by signal propagation over land. DIINS models the delay caused by the atmosphere and signal travel over salt water (the SPF correction). These models are range dependent and very simple. However, the delays due to travel over land are too complicated by geography, geology, etc. to modelled in a simple manner. In practice, the land delays are physically measured along convenient trajectories. Then, various models are employed to estimate the delays away from measurement points (these are called the ASF corrections). DIINS uses polynomial models developed by the Canadian Hydrographic Service: sets of latitude and longitude dependent coefficients were estimated for particular geographic areas. If DIINS is in an area where coefficients are available, it inputs latitude and longitude and receives estimates of the land delay. Unfortunately, these models miss some significant features in the signal delay resulting from propagation over land and improperly model errors at other locations.

All DIINS sea trials have sailed from Esquimalt, B.C. The only Loran-C chain useable in the test area is the Canadian West Coast chain (GRI 5990). The only slave stations tracked during any sea trial have been Y (George, Washington) and Z (Port Hardy, B.C.). These two stations provide good geometry for hyperbolic positioning.

Figure 27 Loran-C Y and Z TD Errors, ST95, Days 129–132

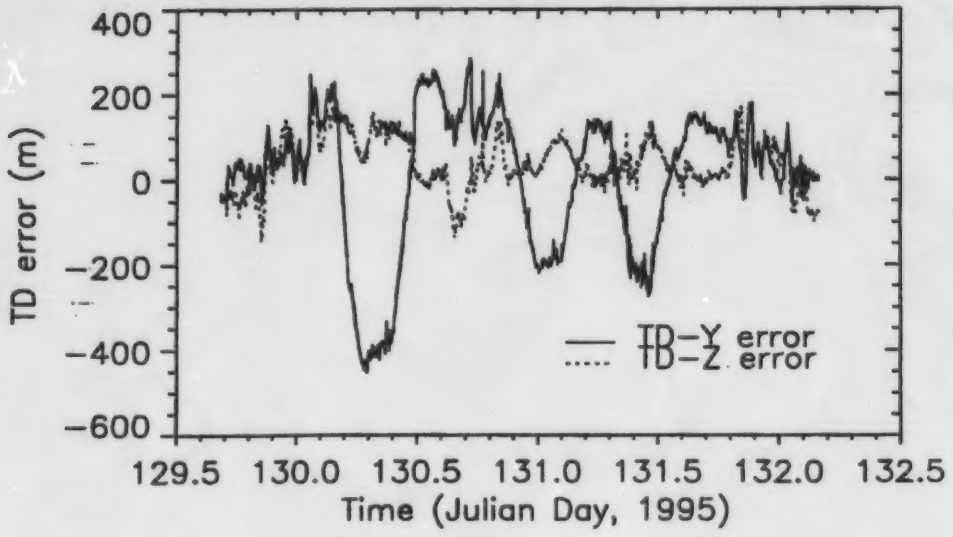

Figure 28 Loran-C Y and Z TD Errors, ST94, Days 322–325

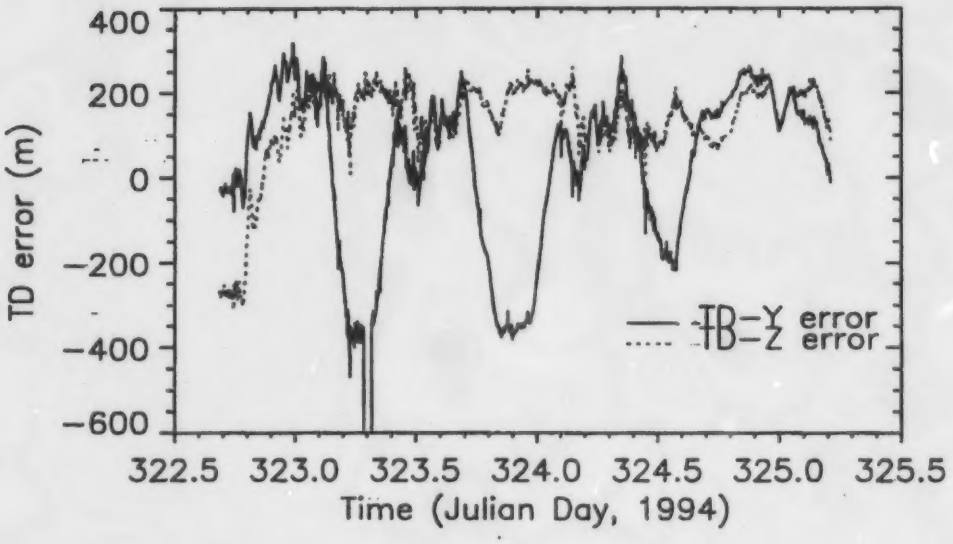

Figure 27 is a time plot of the errors for the Y and Z TDs computed for Days 129–132. The TD errors have been corrected using the SPF and ASF models used in DIINS. Note that the errors in Y are generally larger than the errors in Z. The 5990–Y station at George, Wash. is located well inland: the Loran-C signal from this station to the sea trial test area travels primarily over land. On the other hand, the 5990–Z station at Port Hardy is located at the northwestern tip of Vancouver Island: the signal travels mostly over sea water. If the errors remaining in the plots are largely due to unmodelled ASF errors, the TDs with the longest travel time over land would be expected to have the largest errors. Thus, the data suggests that unmodelled ASF errors are probably the dominant source of systematic error in the corrected TDs. To demonstrate the position dependent, relatively (time) stable nature of these errors, the same TD errors from DIINS Sea Trial 1994 are plotted in Figure 28. The trajectory followed by the Endeavour during this portion of the 1994 sea trial is similar to that followed in 1995.

The section of the common trajectory of most interest in the present analysis is shown in Figure 29. The ship approached the test area from the east arriving at Point 1 then turning south to Point 2. The east to west leg from Point 2 to Point 3 corresponds to the region of large Y TD errors shown at about Day 130.3 in Figure 27 and about Day 323.2 in 28. The large excursion at 323.3, 1994 jumps up to about 6500 metres then down to about 3500 metres and back to 400 meters. At the same time, the signal to noise ratio of the TD-Y signal drops significantly. This excursion is likely caused by (3000 metre) cycle slips. In both sea trials, the ship proceeded to Points 4, 5 and 1. Then, trajectories diverge:

1. In 1994, the ship again sailed from Point 1 to 2 and around the larger, trapezoidal perimeter. The second large Y TD error region on Figure 28 — at about Day 323.8 — occurs as the ship again traverses the leg from Point 2 to 3. Note the similarity in the errors in the first two anomalies (as well as the similarities with the large errors at 130.3, 1995). After proceeding to Points 4, 5 and 1 for a second time, the ship reversed its track, sailed back to Point 5, then turned south to Point 6. The third large Y TD error region — at about Day 324.5 — occurred as the ship traversed the Point 6 to 7 leg. Notice that errors in this third region are similar to the first two but are smaller in magnitude. This leg is near to and parallel to the other area of large errors along the Point 2 to 3 leg. These results would be consistent with the assumption that the large errors are due to unmodelled ASF effects:

the large errors along both legs could be due to the same ASF anomaly.

2. In 1995, after returning to Point 1 for the first time, the ship reversed its path back to Point 5, then turned south to Point 6 and around the inner rectangular path to Points 7, 8 and back to 1. Then, it proceeded around the inner path again. The second and third regions of large TD-Y error (at times 131.0 and 131.4, respectively) occurred as the ship sailed twice from Point 6 to 7. Again, notice the similarities in the error behaviour at these times. Further, notice the similarities with the large errors at 324.5, 1994.

Clearly, there are errors in the Loran-C signals which are position dependent. Given that:

1. the TD-Y signal travels over more land than the TD-Z signal
2. and the fact that TD-Y errors are much larger than TD-Z errors at particular locations,

it seems highly likely that there are unmodelled ASF effects causing the large TD-Y errors.

Figure 29 DIINS Sea Trials' Trajectory

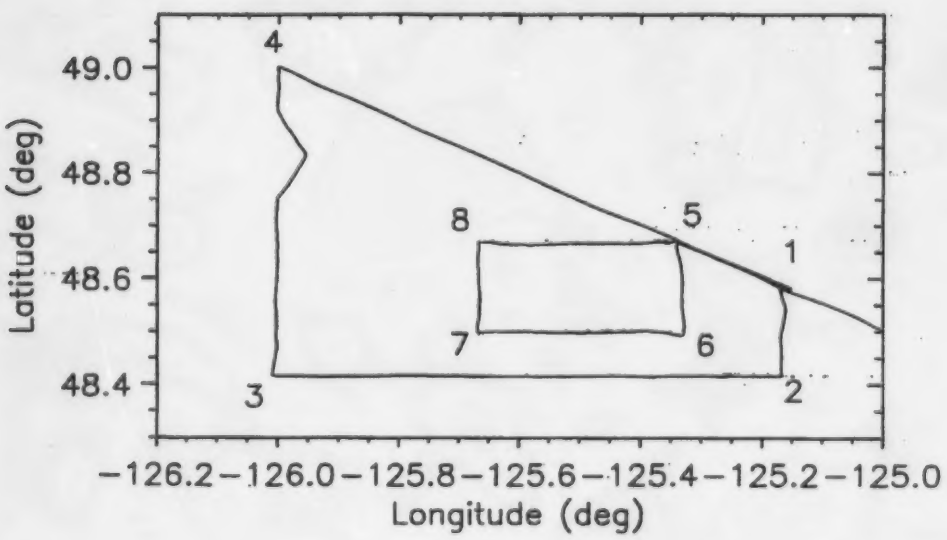

Figure 30 3–D View of Loran-C Y TD Errors, ST94, Days 321–325

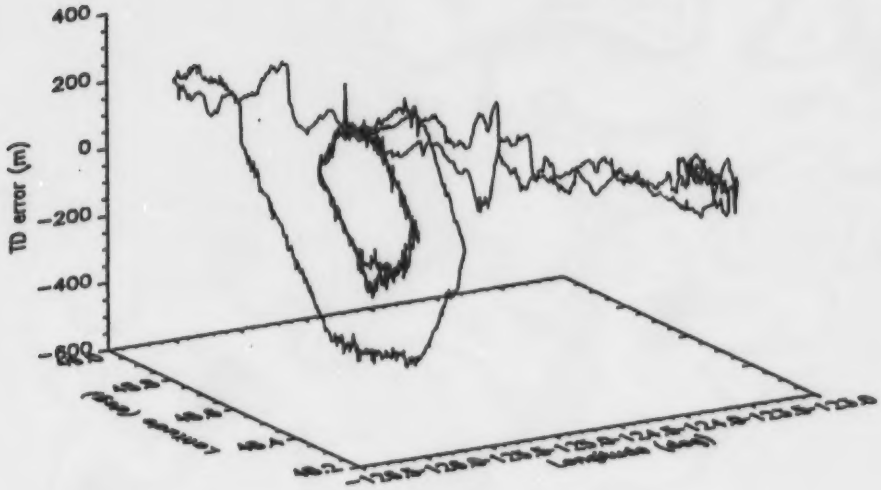

Figure 31 3–D View of Loran-C Y TD Errors, ST95, Days 129–132

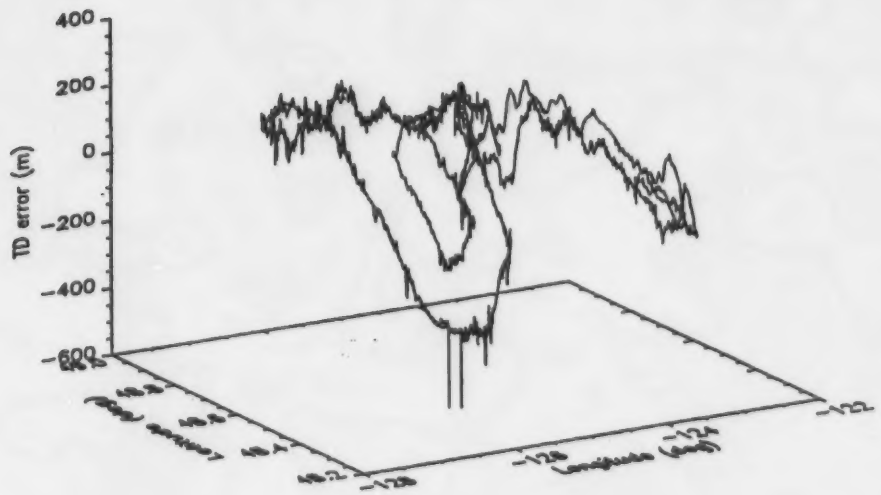

To further demonstrate the position dependent nature of the TD errors, Figures 30 and 31 show three dimensional views of the TD errors from Loran-C slave 5990–Y over the sea trial region. Note the obvious similarities of these two plots.

3.3.2. Filter Models

The Internav LC360 Loran-C time difference errors were each modelled with one state. Since this receiver tracks only one Loran-C chain at a time, and a Loran-C chain will have at most four slave transmitters, four states (one for the TD error of each possible slave) are required for Loran-C. Each is assigned the same model.

As the analysis above explains, there are only TDs from two Loran-C stations available off the west coast of Vancouver Island — 5990–Y and 5990–Z. And, as the analysis showed, the error behaviour of these two TDs was markedly different. This presents two important modelling problems:

1. Are the TD errors observed for these two Loran-C stations representative of the Loran-C system as a whole?
2. Is the error behaviour of either TD more representative of TD errors for other Loran-C stations?

Because of these data limitations, the Loran-C TD error model inherited from DIINS' predecessor, MINS, was used:

1. Measurement noise = 10 metres,
2. First order Gauss-Markov standard deviation = 185 metres,
3. First order Gauss-Markov correlation time = 2.5 hours.

For reference, the models derived from ST93, ST94 and ST95 were developed. The results are presented in Table 17.

Table 17 Loran-C TD Error Models from DIINS ST94 and ST95

Error Model / Sea Trial	1993	1994	1995
TD-Y Meas. Noise (m)	7	8	5
TD-Y G-M Std. Dev. (m)	227	179	164
TD-Y G-M Corr. Time (s)	9700	10,700	8700
TD-Z Meas. Noise (m)	5	7	5
TD-Z G-M Std. Dev.	81	167	74
TD-Z G-M Corr. Time	13,500	>25,000	14,200

Another point of contention arises from the analysis of TD errors in the previous topic. There, the position dependent nature of the main component of the TD errors (the ASF errors) is explained. Here we are attempting to model these (mainly) position dependent errors as time dependent. In this geographical region, the TD-Y errors are larger than the TD-Z errors because of the geography. It is thus not surprising that the autocorrelation functions between TDs are so different while the functions for either TD are so similar from sea trial to sea trial. The models developed from the sea trial data would be appropriate for use in this region. For use elsewhere, they would be inappropriate.

Chapter 4: REAL-TIME FILTER RESULTS: DESCRIPTION AND ANALYSIS

This chapter analyses the results from all seven real time DIINS filters. Any unexpected results are described and where possible explained. All filters shared the following parameters:

1. 2.56 second update rate. This was set to the highest rate the processor would allow by trial and error.
2. The χ^2 test (for soft failure detection) confidence level was set to 99%. The filter failure detection delay was set to 60 seconds.
3. Residual test limits were all set to $n = 5$.
4. The measurement failure detection delay was set to 90% failed over the previous 5 minutes; the measurement failure recovery delay was the same as failure detection (except Filter 1 — see below).
5. All filters utilized 7 system states (2 position, 2 velocity and 3 attitude) and 5 INS sensor error states (2 accelerometers biases and 3 gyro biases).
6. All filters used 2 states and 2 measurements to model the synthetic velocity measurements used to correct for the internal MK29 velocity damping.
7. All sensor error models are the same in all filters (except Filter 1 below).

4.1: Filter 1 (INS 1 with INS 2 Position Measurements)

This filter is used exclusively for Failure Detection, Isolation and Reconfiguration (FDIR) and provides no useful navigation information. It used the fore MK29(1) as the primary unit and processed aft MK29(2) low noise latitude and longitude measurements. INS sensor error variances were set to twice the values used in other filters to reflect the fact that the filter is estimating INS1/INS2 error differences. Only the 12 INS states and 2 synthetic velocity error states listed above were used in this filter. The measurement failure detection delay was set to 100% failed over the previous 100 seconds; the measurement failure recovery delay was set to 100% passed over the previous 50 seconds.

This filter worked well during the sea trial: there were no measurement failures or unusual behaviour.

4.2: Filter 2 (INS 1 with TASMAN P(Y) GPS Position and LC360 Loran-C Time Difference Measurements)

This is one of the unifilters. As such, it is also one of the largest:

1. It processed latitude and longitude measurements from the Trimble TASMAN P(Y) code GPS receiver as well as Loran-C time differences.
2. Along with the basic 14 states, there are 2 additional GPS position error states, 1 additional GPS time delay state and 4 additional Loran-C time difference states, for a total of 21 states in the filter.

The number of Loran-C TD measurements processed at any particular epoch depends on the number of TDs measured by the receiver: this can vary from 0 to 4. During ST95, Loran-C data was collected from only 2 secondary stations (Y and Z) in one chain (Canadian West Coast, GRI 5990), ie. only 2 TD measurements were processed per epoch. However, in the interest of completeness, all four secondary measurements and their associated states were activated.

4.2.1. Anomaly 1

Anomaly Observed

DIINS position errors and standard deviations are larger than expected (100–200 metres) from the start of Run 1 (74,276,507 seconds or 129:16:22) until about 74,372,000 seconds (130:19). Occasionally, the standard deviations drop down to expected values (around 10 metres). Figure 32 illustrates this anomalous behaviour. Similar behaviour is also observed during Run 4 (Day 136).

Figure 32 Filter2: Position Error Anomalies

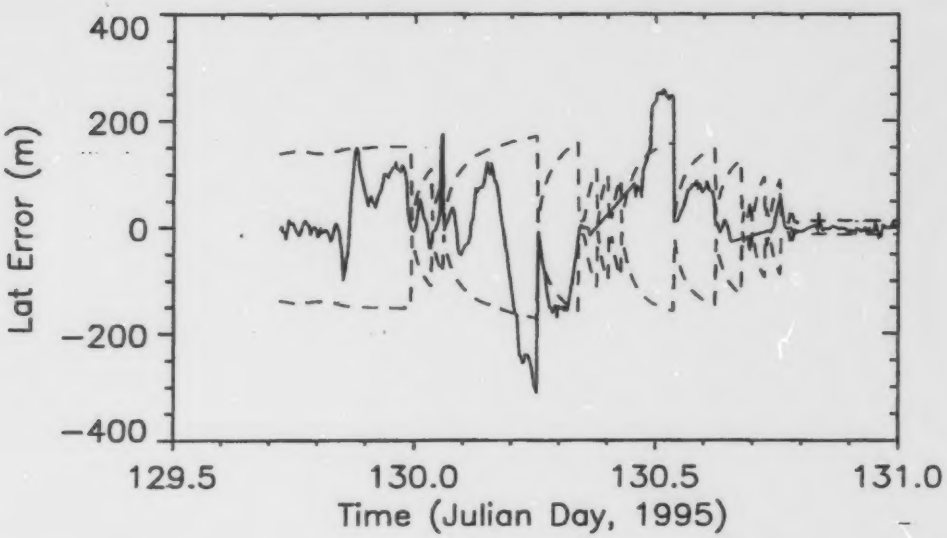

Cause of the Anomaly

The DIINS' FDI log for Run 1 contains the warning "FDI: Sensor GPS1_SNSR NO measurements being processed at 74823872.772" (seconds or 128.93 days). Examination of the sea trial log books and discussions with the personnel operating DIINS during the trial revealed no evidence of GPS receiver problems. Post-processing of the troublesome data did not produce any GPS losses: GPS measurements were continuously processed. The problem with the real time system originated with small irregularities in the arrival times of the GPS data records which were triggering a fault detection mechanism in the real time Data Collection Task. This resulted in almost all GPS records being sent to the filter processing software with status flags set to bad. The filters using GPS did not form measurements with the "bad" GPS data. However, the "unifilters" (2 and 3) were processing Loran-C time differences. The outputs from Filters 2 and 3 reflect the presence of those Loran-C measurements.

Occasionally, one or more GPS records were flagged good. These records were used to form measurements which temporarily reduced the position errors and standard deviations. Then, the time would deviate again, the status flag was set to bad once again, and the GPS measurements were disabled. Since these

periods of good status were typically short, there was little opportunity for the GPS measurements to model INS errors and the filter errors quickly rebounded to the Loran-C range.

In most cases, the FDI log files output by DIINS do not give any indication that these measurements are being processed. The DIINS FDIR processor active during ST95 did not do anything except notify the operator when good aiding sensor data was lost, ie. GPS measurements were not disabled when the "FDI: Sensor GPS1_SNSR NO measurements being processed . . . " message was sent. So even one GPS record with a good status message would be processed. Note that this message was only sent after a span of time equal to 'ChiSquared.FailureTriggerTime' had passed with no GPS measurements processed. For ST95, 'ChiSquared.FailureTriggerTime' was set to 60 seconds. Similarly, the message indicating that GPS measurement processing had resumed ("FDI: Sensor GPS1_SNSR measurements being PROCESSED at 74334135.927" — 130.35 days) was sent only after a span of time equal to 'ChiSquared.FailureTriggerTime' had passed with no interruption in GPS measurement processing. This did occur on four separate occasions. However, in most cases, GPS status recovered for a length of time which was less than the required 60 seconds: the GPS measurements were processed but no FDI recovery message was signalled.

Additional similar FDI messages were logged as follows:

1. Run 1 (Days 128–132) —
   a. "FDI: Sensor GPS1_SNSR NO measurements being processed at 74334432.887" (130.35 days).
   b. "FDI: Sensor GPS1_SNSR measurements being PROCESSED at 74376219.758" (130.84 days).
2. Run 2 (Day 133) —
   a. "FDI: Sensor GPS1_SNSR NO measurements being processed at 74660814.892" (134.13 days).
   b. "FDI: Sensor GPS1_SNSR measurements being PROCESSED at 74660904.492" (134.13 days).
   c. "FDI: Sensor GPS1_SNSR NO measurements being processed at 74661806.893" (134.14 days).

d. "FDI: Sensor GPS1_SNSR measurements being PROCESSED at 74661895.211" (134.14 days).

3. Run 3 (Days 134–135) —

   a. "FDI: Sensor GPS1_SNSR NO measurements being processed at 74766992.146" (135.36 days).

   b. "FDI: Sensor GPS1_SNSR measurements being PROCESSED at 74802796.297" (135.77 days).

4. Run 4 (Day 136) —

   a. "FDI: Sensor GPS1_SNSR NO measurements being processed at 74823872.772" (136.02 days).

   b. No recovery was logged.

Fix of Anomalous Condition

The fault detection mechanism in the real time software which caused the GPS data to be flagged 'Bad' has been modified. This algorithm was designed to detect stale GPS records. The time limit used previously was too small to handle the data coming from the TASMAN P(Y) GPS receiver. The problem was fixed simply by increasing the detection time threshold. Lab tests using the modified real time software with a TASMAN receiver and one Sperry MK29 INS have shown that the problem has been solved: only a very small percentage of the GPS records are now being flagged.

4.2.2. Anomaly 2

Anomaly Observed At about time 74,760,000 seconds (135.28) the behaviour of the DIINS position standard deviations changes from "smooth" to "noisy", and, at the same time drops from about 14 metres to about 9 metres. There is no change in the GPS supplied position variances, which should be the main factor in the DIINS position error estimate. Closer examination of the DIINS

state estimates shows other states behaving in a similar way. Figure 33 illustrates this anomalous behaviour.

Figure 33 Filter2: Position Standard Deviation Anomaly

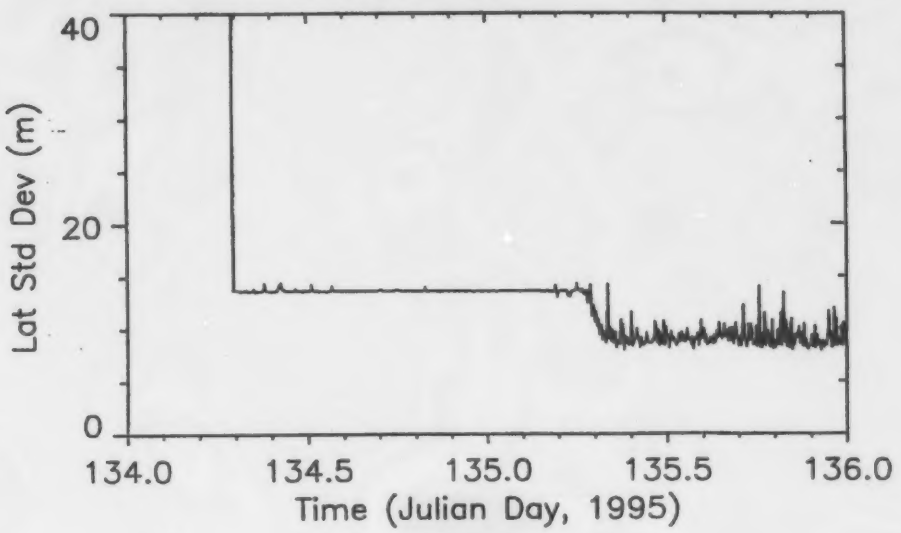

Cause of the Anomaly There is no indication of any problems at this time from any of the log files output as DIINS was running. These logs report filter generated messages (χ^2 failures, measurement failures, changes in velocity damping status, etc.) and FDI messages (sensor failures, changes in the filter selected for output). However, the FDI log does report that "FDI: Sensor GPS1_SNSR NO measurements being processed at 74766992" 1.5 to 2 hours after the start of the anomaly. Notice that the DIINS estimated position standard deviations do not reflect any loss of GPS data except for a small spike in position standard deviations: this is an indication that GPS may have been off just long enough to trigger the FDI message. The fact that the FDI message signalling resumption of GPS processing was not displayed indicates that there is no 60 second span with consistently good GPS status subsequent to the start of this anomaly. This further indicates that GPS data was unreliable during ST95 real time operation and a possible cause of the anomaly. This is discussed in more detail under Topic 5.2.1.

Again, this anomaly cannot be reproduced in post-processing or lab trials. Some further observations suggesting that this anomaly has its roots in a problem with GPS:

1. The GPS position state standard deviations behave just like the system (INS) positions. The GPS states are likely driving the system states.
2. The GPS time delay state estimate becomes less "noisy" and moves from a mean value of about 0.003 seconds to about 0.000 seconds. In addition, this state's standard deviation changes from a value oscillating about 0.0498 seconds to a fairly constant value of about 0.050 seconds, the predefined steady state value. These factors suggest that the GPS time delay state has become unobservable.
3. This behaviour is observed in every filter using GPS and absent from all other filters; again pointing to a problem with GPS.

Figure 34 Filter 2 Latitude Residual with Failure Bounds (5 sigma)

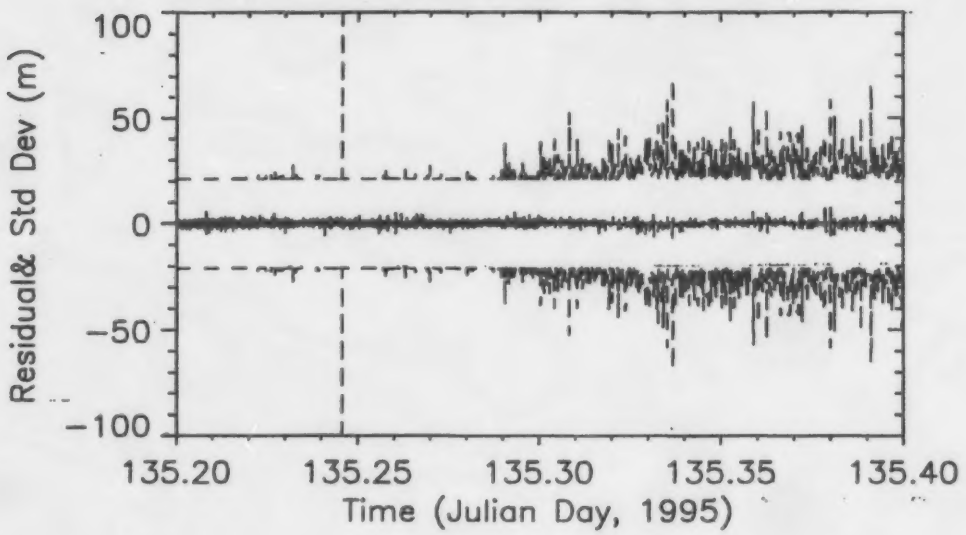

Figure 35 Filter 2 Latitude Residual — Before Anomaly

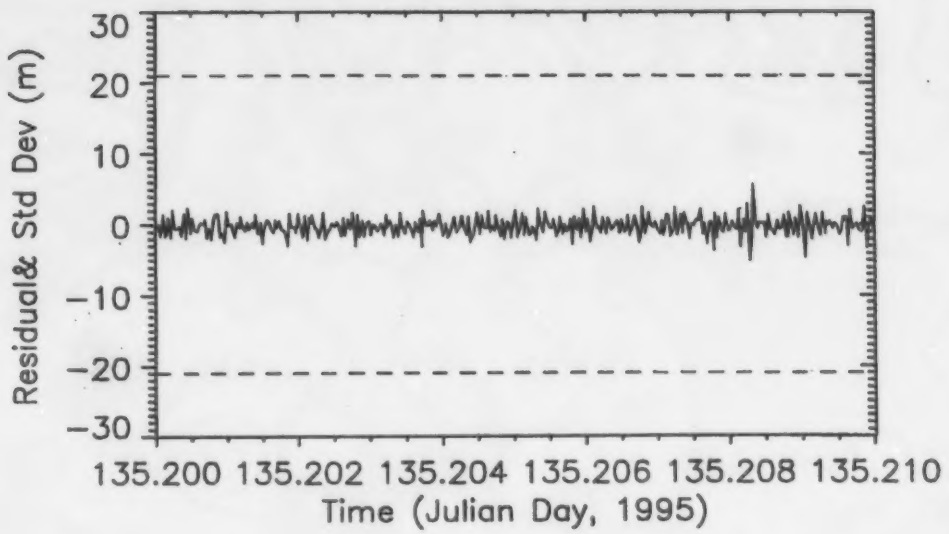

Figure 36 Filter 2 Latitude Residual — After Anomaly

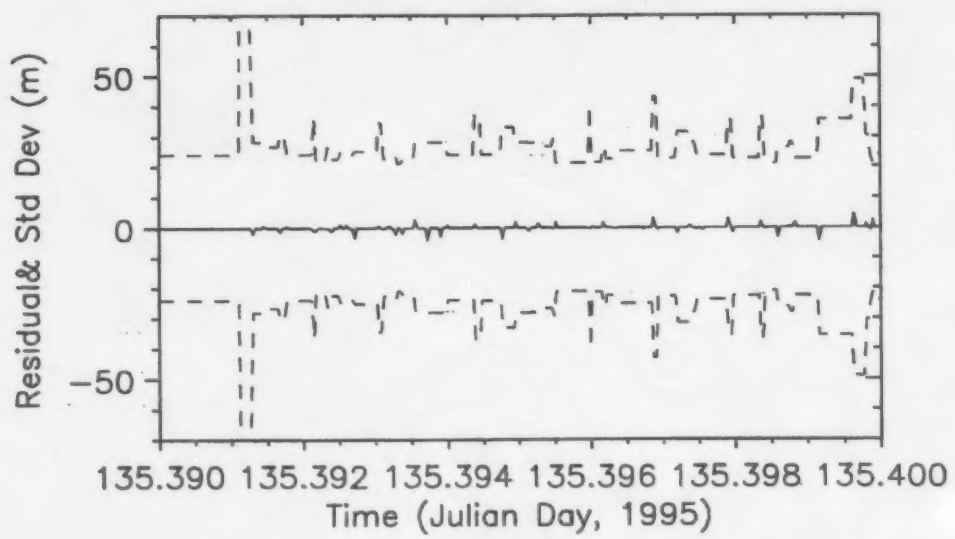

In DIINS, residuals are computed only when good data from the primary INS and the aiding sensor are available. Residuals which are not computed at any particular epoch are set to zero; their standard deviation estimates are not changed. The residuals and their estimated standard deviations were output and stored during ST95. An examination of the residuals from this filter confirmed that measurement density decreases at about the same time as the anomaly started. Figure 34 shows the latitude residual with its 5–sigma failure bounds in the vicinity of the start of the anomaly. Figures 35 and 36 are expanded views before and after the start of the anomaly, respectively.

Notice all the zero valued residuals and coincident constant standard deviation estimates in Figure 36. This indicates that residuals were not computed at these times. Thus, the start of this anomaly starts at the same time as significant numbers of GPS data records start being flagged bad. These two observations are almost certainly related. However, how can the DIINS' position standard deviation estimates be reduced in magnitude when fewer measurements are being processed? The answer to this question lies in an algorithm error which has since been corrected.

When position accuracy estimates are available from a GPS receiver, DIINS uses them to modify the process noise model for GPS positions. For ST95, the TASMAN GPS receiver output a kind of distance RMS (dRMS) — the RMS of latitude and longitude standard deviation. In DIINS, the TASMAN latitude and longitude standard deviations are assumed to be equal and are each set to the dRMS times the square root of 2. The resulting variances are assumed to be comprised of a constant measurement noise component plus a process error component that changes as the dRMS changes. A detailed discussion of this process is beyond the scope of this report. In brief, the GPS position errors are modelled as first order Gauss-Markov processes. The (continuous) spectral density matrix of such errors is completely defined by the steady state variance and the correlation time. The discrete Kalman filter used in DIINS requires the (discrete) variance of each GPS position error state. This is a function of the spectral density and the propagation time of each state. The algorithm to use the TASMAN dRMSs by separating them into measurement and process components originally assigned the process component to the discrete variance. This effectively treated the process error as white noise rather than the required Gauss-Markov process. It is more appropriate to set the process component of the receiver error to the steady state

variance of the spectral density. This is then used to compute the discrete variance. The incorrect model was used for ST95 but has since been corrected.

So how does all this relate to the anomaly under investigation? Recall that the DIINS position variances decreased as the frequency of GPS measurements decreased. Fewer measurements should have produced larger variances. However, when no good data is available from the GPS receiver, the process error remodelling scheme is not activated. This is the case for the old algorithm and the new. However, here the similarity ends. Recall that the original algorithm set the process error component of the receiver's estimate of the position accuracy to the discrete variances and did not change the continuous spectral densities. At every epoch, the variance is zeroed and recomputed using the spectral density. Thus, when good GPS was not available, the GPS state variances were set using the a priori models. Because the GPS receiver's dRMSs were used incorrectly, the a priori model produced DIINS position variances which were consistently lower than those computed using the dRMSs.

Fix of Anomalous Condition This anomaly required the following fixes:

1. Modify the fault detection mechanism used in the real time software to flag bad GPS data (as discussed in 4.2.1). The necessary changes have been made and lab tests indicate they are working properly.
2. Correct the algorithm which uses position accuracy estimates from the GPS receiver to modify the process error variance of the GPS position states. When the process error component of the receiver's accuracy estimate is correctly assigned to the continuous spectral density (instead of the discrete variance), a loss of GPS data leads to a fixed spectral density model. At present, the model is fixed at the last values computed using the receiver information. Alternatively, it could be reset to the a priori model.

4.3: Filter 3 (INS 2 with TASMAN P(Y) GPS Positions and LC360 Loran-C Time Differences)

Recall that Filter 2 and Filter 3 differ only in the INS being filtered: the configuration of Filter 3 is exactly the same as Filter 2 (as described in the preceding section). In summary, GPS latitude and longitude, and 4 Loran-C TD measurements were activated (although only 2 were ever processed); a total of 21 states were modelled.

All the GPS-based anomalies reported above for Filter 2 also apply to Filter 3. Additional anomalies not observed with Filter 2 are reported below.

4.3.1. Anomaly 1

Anomaly Observed Longitude errors for parts of Runs 2 and 3 (Days 133–135) are worse than DIINS expects, worse than longitude errors from Filter 2, and worse than latitude errors. Figures 37 and 38 illustrate this. Note the large longitude errors at times 133.8, 134.2, 134.4 and 134.8 days.

Cause of the Anomaly Because this anomaly was not observed in Filter 2, based on MK29–1, and because everything else (software, aiding sensors, etc.) is the same as Filter 2, it seems likely that these larger than expected longitude errors are related to the performance of MK29–2. First of all, let's take a look at the position errors from MK29–2. Latitude and longitude errors (vs the Trimble TASMAN P(Y) code GPS receiver), along with the internal SKOR filter estimates of standard deviations, are shown in Figures 39 and 40. Once again, the latitude errors are smaller than longitude errors and are mostly within the 1–sigma bounds; the longitude errors are larger than latitude errors and are well outside the 1–sigma bounds much of the time, especially at the times when the DIINS errors were larger than expected. Also, note the distinctive Schuler—like error behaviour in the MK29–2 outputs at the times of larger than expected longitude errors. There are two speed log based velocity damping filters used in the MK29s (one north and one east) which are supposed to remove the Schuler oscillations. However, in some circumstances, these damping filters will be deactivated. Figure 41 shows the status of the damping loops. In most cases, times when the damping loops are deactivated coincide with the large longitude errors. The exception is at the start of Run 3 (Days 134–135) — at about 134.4 days. Here, the damping loops are active but

1. the DIINS longitude errors are large,
2. and, the MK29–2 longitude errors display Schuler-like oscillations (even though they are mostly within the SKOR 1–sigma bounds).

Another area of interest occurs at time 134.20 where the DIINS longitude error spikes sharply up to 100 metres. This can be explained by looking at Figure 42. This figure shows the filter damping "Mode" selected by the operator:

1. It can be set to AUTOmatic mode where the INS will decide whether the damping loops should be activated or not based on the difference between the speed log reference and the INS velocities.
2. Or, it can be set to MANual Damped or Undamped mode. In this mode, the operator decides if the damping loops should be activated. The MANU mode would be selected if the operator knew that the speed log was unreliable. MAND can be selected when AUTO mode has deactivated the loops but the operator feels that the loops can be safely reactivated. Proper operation of the speed log should be confirmed before MAND is selected. Further, since the speed log reference is usually measuring speed through the water and not speed over ground, it will be unreliable during manoeuvres. Normally AUTO mode will deactivate the damping loops during manoeuvres. Figures 43 and 44 show the trajectory (latitude vs longitude) and longitude (vs time). Good examples of the AUTO damping mode switching to undamped during manoeuvres can be seen from time 134.5 to 134.7: each switch to damping off (Figure 41) results from a turn (Figure 44).

At the large longitude error spike at time 134.20, damping is in the MAND mode. At the same time, the ship executes a 90° turn to port (from north to west). Since the MAND damping mode is set, the damping loops are forced to accept the speed log velocities even when those velocities are unreliable through the turn. This leads to errors in the INS velocities which leads to errors in the INS positions.

Figure 37 Filter 3 Latitude Errors with DIINS Standard Deviations, Day 134

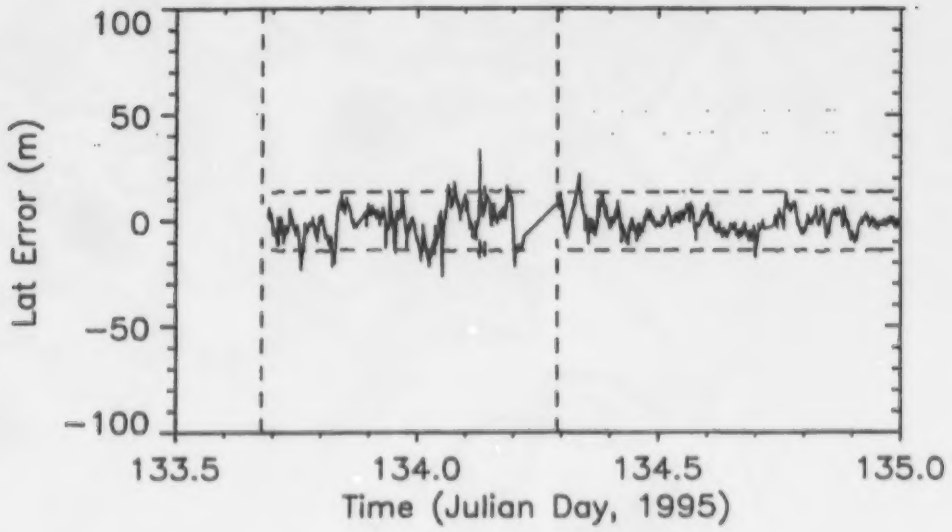

Figure 38 Filter 3 Longitude Errors with DIINS Standard Deviations, Day 134

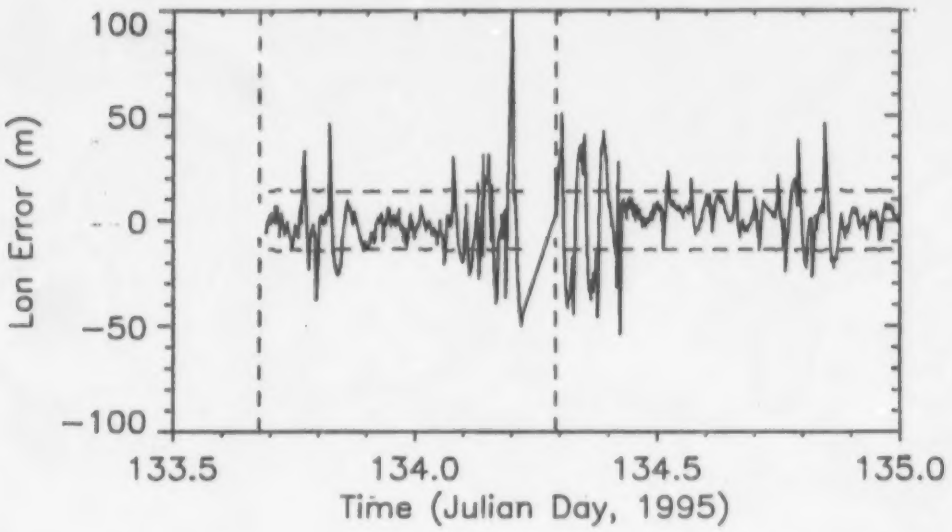

Figure 39 MK29-2 Latitude Error vs GPS with SKOR Estimated Standard Deviations, Day 133-134

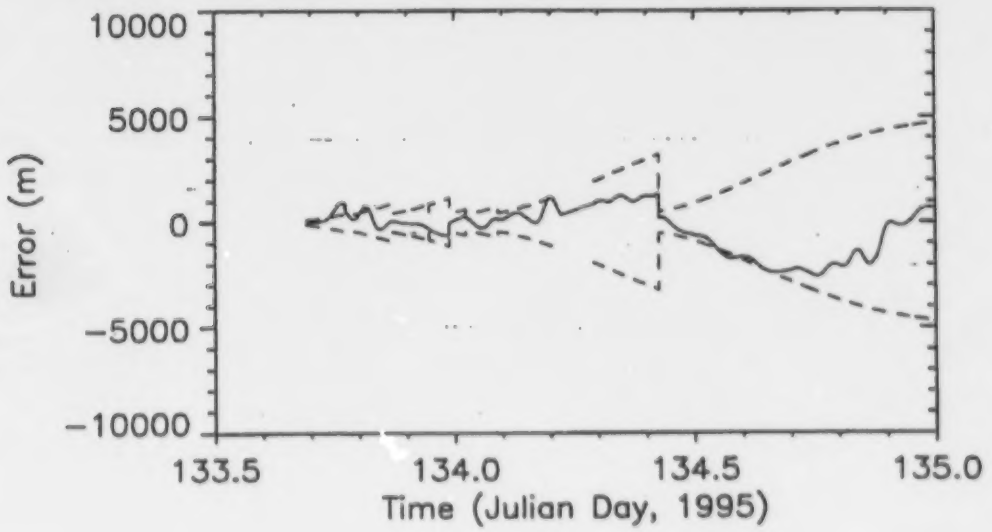

Figure 40 MK29-2 Longitude Error vs GPS with SKOR Estimated Standard Deviations, Day 133-134

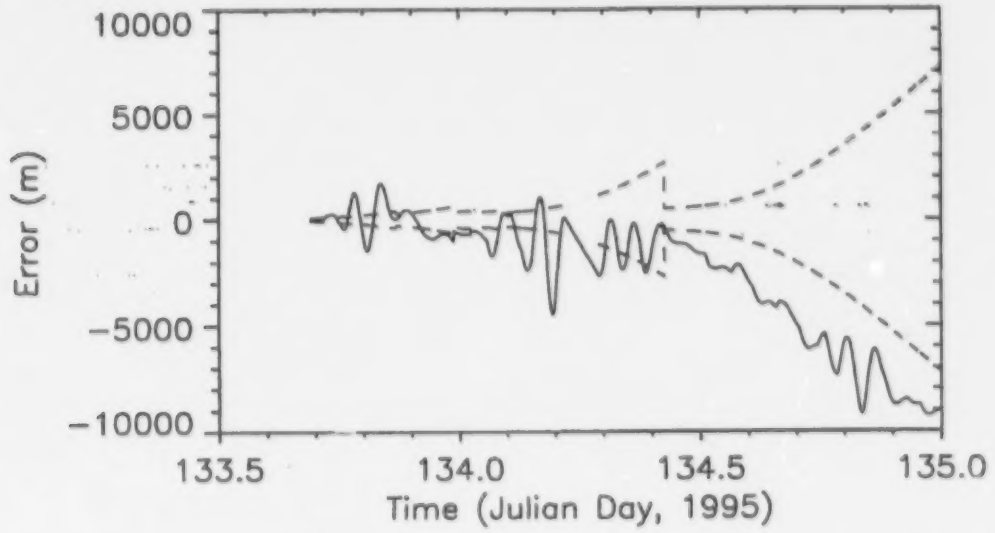

The same reasoning can be used to explain the large DIINS longitude errors (and oscillating INS longitude errors) around time 134.4 days:

1. The MAND damping mode is still set (AUTO was selected for a few seconds at 134.40 but was not left on).
2. A number of manoeuvres were executed between the start of Run 4 (134.29) and the operator's reselection of AUTO damping mode (at 134.47):

   a. some small turns at the start of the run;
   b. a 90° turn to port (from north to west) at 134.30;
   c. a 90° turn to starboard (from west to north) at 134.325;
   d. a 90° turn to port (from north to west) at 134.35;
   e. a 90° turn to starboard (from west to north) at 134.375;
   f. more brief manoeuvres around 134.42

   In fact, these manoeuvres occur at a regular rate which is close to the 84 minute Schuler rate. This may result in some extra "boosting" of the Schuler effects.

The forced use of speed log damping through manoeuvres seems to induce the INS to oscillate in Schuler-like manner.

While looking at the damping performance, a peculiar phenomenon was chanced upon. Figure 45 shows the MK29–2 east velocity error (vs GPS) and the MK29–2 estimate of the east ocean current. Notice that the behaviour is very similar. It appears that the east velocity errors are being driven by the east ocean current estimate. For the north velocity of MK29–2 and both velocities for MK29–1 the effect is less apparent. However, since the errors are smaller in these cases, it may be that the correlation is just less obvious. There is no information available which explains how the ocean currents are computed. A close-up look at the data revealed that

1. ocean currents are estimated even when the damping loops have been deactivated;
2. ocean current estimates change in as little as 1.28 seconds.

The first point suggests that the ocean currents are not computed in the damping loops. They may be estimated in the SKOR filter: Reference [7] states that "*After every Kalman filter update, . . . north and east velocity resets are applied; apparently, these resets are somehow computed by the Kalman filter . . .* ". Further, an early proposal written by Sperry Marine (reference [8], page 3–1) states:

"Speed log bias errors do not influence accuracy performance during straight sailing because of the use of third-order vertical loop damping. . . . The velocity reference error (that is, the difference between ground speed and water speed) may then be considered to consist only of ocean currents. . . . random ocean currents are modelled by an exponential correlation function with a correlation distance of 15 nautical miles and an RMS value of 0.6 knot." There is no known 1.28 second loop in the SKOR processing but the possibility cannot be excluded. Also, reference [9], page 4, explicitly states that there are "Inertial Velocity North" and "East" states in the SKOR filter but does not explain how they are mechanized.

Notice that the SKOR filter's accuracy estimates in Figures 39 and 40 do not demonstrate any Schuler oscillation even when the velocity damping loops are inactive and the INS errors are oscillating. This indicates that the SKOR filter assumes that the Schuler oscillations have been removed via the damping loops, even when the damping loops have been deactivated. Thus, the operator should NOT rely on the SKOR standard deviation estimates when damping is off.

It is not clear why longitude is more affected by these circumstances than latitude. There are independent damping loops for north and east velocities. Is it possible for the north velocities to be damped while the east velocities are not? Figure 46 shows the velocity damping corrections from both the north and east loops along with the damping flag. Whenever the damping flag is TRUE, both loops calculate INS velocity corrections. Similarly, whenever the flag is FALSE, neither loop computes corrections. So it seems clear that both damping loops are activated and deactivated together — this is not the source of difference between latitude and longitude performance.

Figure 41 MK29–2 Inertial Velocity Damping ON Flag, Day 133-134

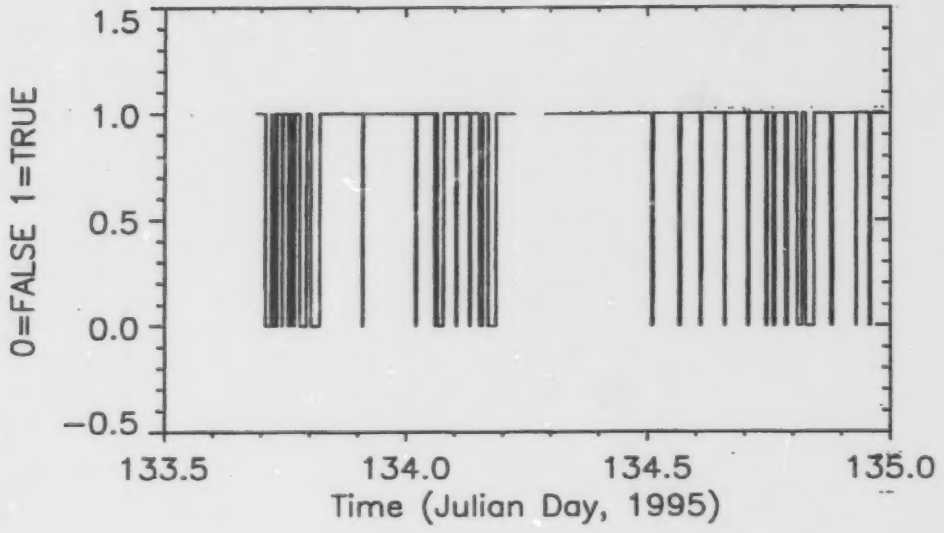

Figure 42 MK29–2 Velocity Damping Mode Flag, Day 133-134

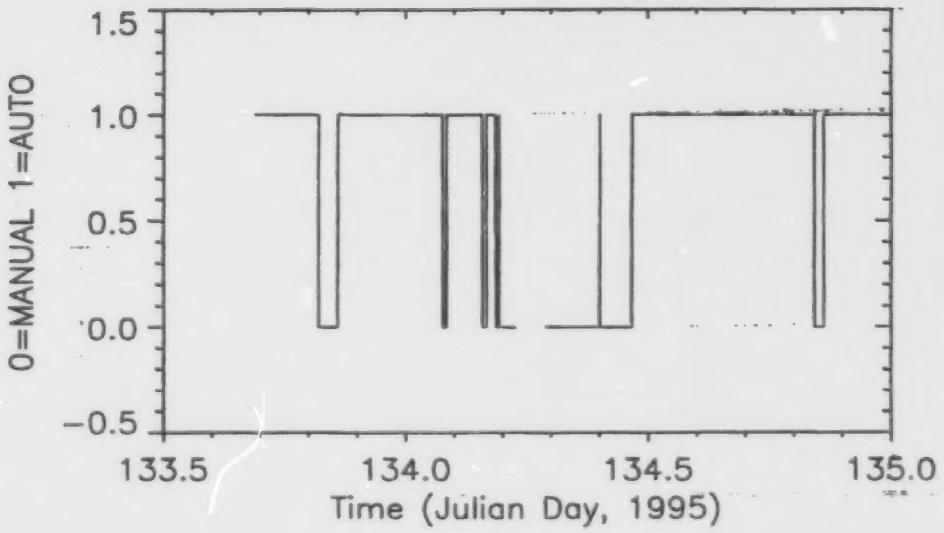

Figure 43 DIINS Trajectory, Day 133-134

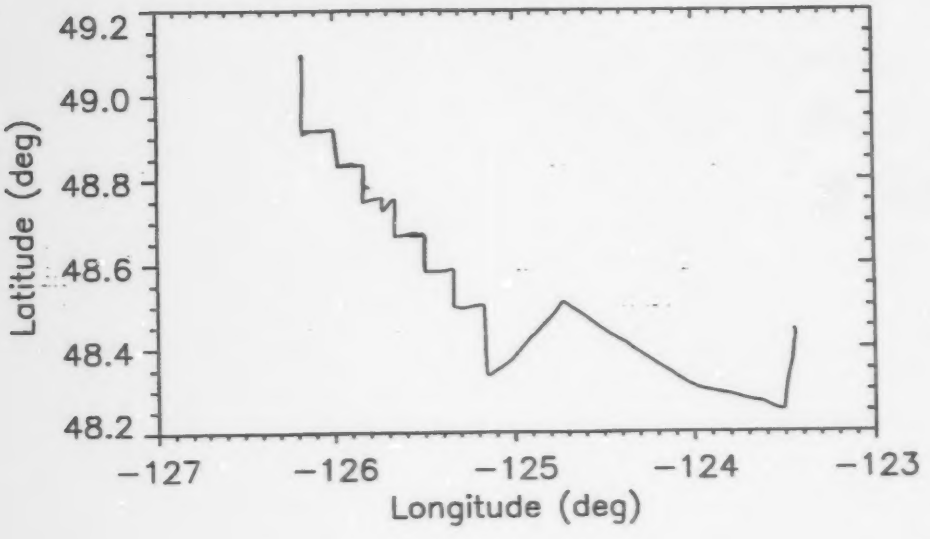

Figure 44 DIINS Longitude, Day 133-134

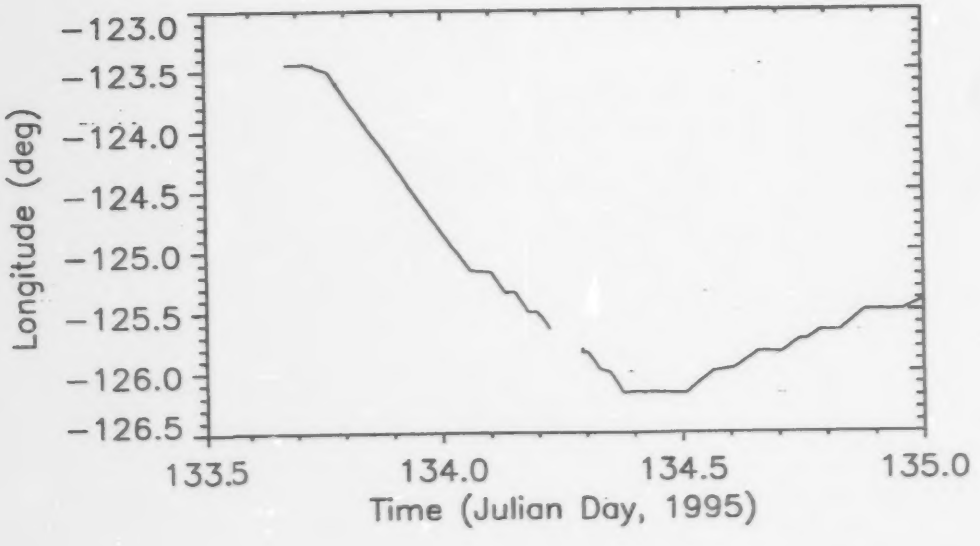

Figure 45 MK29–2 East Velocity Error vs East Ocean Current Estimate

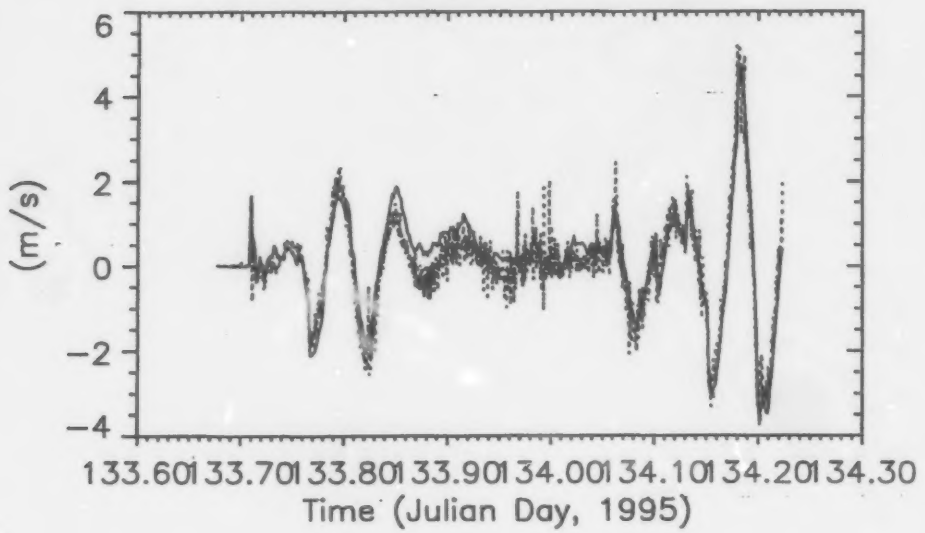

Figure 46 MK29–2 North and East Velocity Corrections (with Damp Flag), Day 133

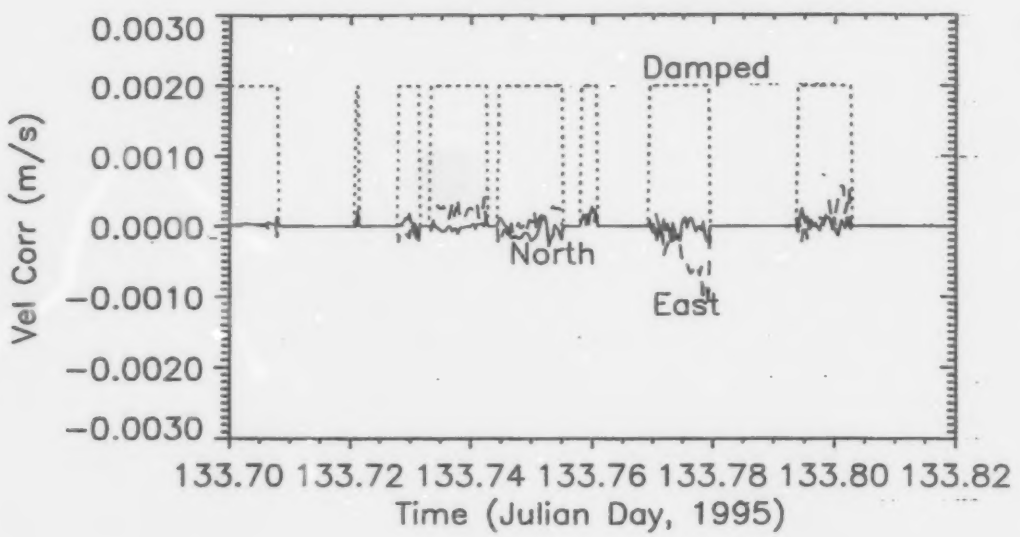

Figure 47 MK29–2 Performance Degraded Status Flag, Day 133–134

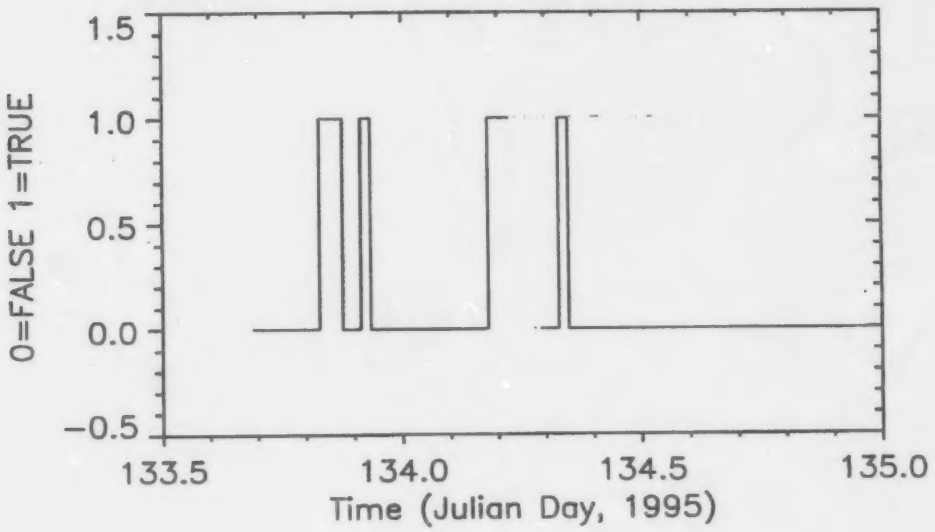

Figure 48 MK29–2 Inertial Well Settled Status Flag, Day 133-134

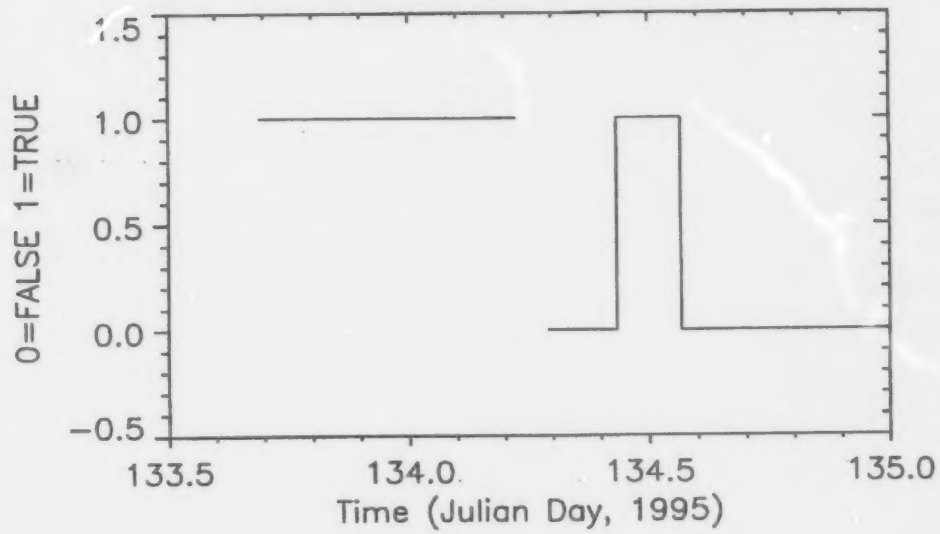

Other possible causes of the latitude / longitude error differences include trajectory effects and INS sensor errors:

1. The differences may be related to this particular trajectory. Most of Run 2 (Day 133) was spent travelling to the test site — primarily in a westerly direction. All of Run 3 (Days 134 to 135) was spent travelling up and down the "staircase". The times of larger longitude error do not seem to be correlated with any trajectory feature.
2. Since the MK29s are stabilized platform INSs with sensitive axes maintained in north-south, east-west orientation, the same accelerometer is always measuring north-south or east-west accelerations. Similarly, a constant platform tilt will produce errors in the (constant) direction of tilt. The DIINS accelerometer states are about the same size in north-south and east-west axes; but the north-south axis is slightly more observable according to the filter. The tilt states from DIINS are converted to roll, pitch, heading format so they cannot be directly related to any particular direction. However, it seems that INS sensor error is the more likely cause of the latitude / longitude discrepancies.

Whatever the cause of these anomalies, can they be identified via MK29 status messages? The damping mode and damping active status is available from both the MK29 display and the digital output records. These have been discussed above. When the damping loops are deactivated, the SKOR filter position accuracy estimates are more likely to be unreliable — they seem to be computed with no knowledge of the damping status. Two other status flags may provide additional insight into MK29 performance: they are the Performance Degraded flag and the Inertial Well Settled flag. Apparently, neither flag is available anywhere but the output data stream: there is no known indication from the MK29 control and display panel. Figures 23 to 24 show the settings over the entire sea trial. Let's look at Runs 2 and 3 in detail to try to identify any patterns.

The flags are displayed in Figures 47 and 48, and summarized below.

1. The Performance Degraded flag changed at the following times (under the following conditions):
    a. TRUE at 133.83 (630 seconds after MAND was set); FALSE at 133.88 (800 seconds after a position fix, 1370 seconds after AUTOD).
    b. TRUE at 133.91 (240 seconds after damping enabled, 400 seconds after damping disabled); FALSE at 133.94 (nothing to report).

c. TRUE at 134.18 — Run 3 (820 seconds after damping disabled); FALSE between Runs 3 and 4 (no data).
d. TRUE at 134.33 — Run 4 (nothing to report); FALSE at 134.35 (nothing to report).

Like the problems discussed in Topic 3.2.3, it is not clear why the flag was set in any of these circumstances. The Sperry manuals say that the flag is only set when

- the INS has run for more than 10 hours without a position fix;
- or, the internal (SKOR) heading standard deviation estimate exceeds 5.0 arc minutes.

Neither limit is even approached in any of the above examples: perhaps the Performance Degraded flag is also, in some way, dependent on velocity damping performance.

2. The Inertial Well Settled flag was set to TRUE (the no fail case) throughout Run 2. In Run 3, it changed at the following times (under the following conditions):

   a. FALSE (the failure case) between Runs 2 and 3 (while data was not being recorded), TRUE at 134.43 (650 seconds after a position fix);
   b. FALSE at 134.57 (170 seconds after damping disabled), TRUE at 135.25 (220 seconds after a position fix);
   c. FALSE at 135.35 (330 seconds after damping enabled, 670 seconds after damping disabled, etc.), did not recover;

   There is no clear explanation of the meaning or consequences of this flag's failure in the Sperry manuals. From looking at this data, there is no clear cause for the Inertial Well Settled condition to fail, but it seems clear that it is most likely to recover after a position fix (this is confirmed by Run 1 data).

Fix of Anomalous Condition The only clear signal from all of this is operational: a MK29 INS should not be forced in Manual velocity damping mode when it is executing any manoeuvres. Also, the operator should be aware that the SKOR position standard deviations (σ_{LAT} and σ_{LON} on the display) may be less reliable when velocity damping is disabled — the SKOR filter seems to assume that damping is always enabled.

4.4: Filter 4 (INS 1 with TASMAN P(Y) GPS Positions)

Filter 4 belongs to the second level of filters. It uses GPS position measurements and states but does not include any Loran-C measurements or states. This filter is smaller than the unifilters:

1. It processed latitude and longitude measurements from the Trimble TASMAN P(Y) code GPS receiver.
2. Along with the basic 14 states, there are 2 additional GPS position error states and 1 additional GPS time delay state, for a total of 17 states in the filter.

Real-time anomalies for Filter 4 are similar to those observed for Filter 2:

1. GPS measurements are sporadic at the start and end of the sea trial. Since this filter does not have any Loran-C measurements, the DIINS' position errors and standard deviations are larger than those for Filter 2 during the periods where GPS measurements are not available.
2. The behaviour of the DIINS position standard deviations changes at about time 135.28 days.

No additional insight into the causes of these anomalies presented itself.

4.5: Filter 5 (INS 2 with TASMAN P(Y) GPS Positions)

Recall that Filter 4 and Filter 5 differ only in the INS being filtered: the configuration of Filter 5 is exactly the same as Filter 4 (as described in the preceding section). In summary, GPS latitude and longitude measurements were activated and a total of 17 states were modelled.

All the anomalies reported above for Filter 3 also apply to Filter 5 (none of the Filter 3 anomalies are believed to be Loran-C based). No additional anomalies were observed with Filter 5.

4.6: Filter 6 (INS 1 with LC360 Loran-C Time Differences)

Filter 6 also belongs to the second level of filters. It uses up to 4 Loran-C TD measurements and 4 states but does not include any GPS measurements or states. This filter is smaller than the unifilters but larger than filters 4 and 5:

1. It processed TD measurements from the Internav LC360 Loran-C receiver.

2. Along with the basic 14 states, there are 4 additional Loran-C TD states for a total of 18 states in the filter.

Real-time anomalies which were observed for Filter 6 are described below.

4.6.1. Anomaly 1

Anomaly Observed Latitude errors are larger than the DIINS' expected standard deviations, but the latitude standard deviations (~70 metres) are smaller than the longitude standard deviations (~175 metres). The 175 metres longitude standard deviations are more reasonable.

Cause of the Anomaly To try to discover the cause of the latitude standard deviation anomaly, the data was run post-mission with the same filter models. The result was quite surprising:

1. The latitude and longitude standard deviations are roughly the same size (latitude ranges from 130–190 metres, longitude ranges from 165–210 metres).
2. The post-mission DIINS position errors are much worse than the real-time results.

After much analysis, it was discovered that differences in the way that the MK29 position fixes were processed in the two runs were causing the different error behaviour. In particular, in the real-time software used for ST95, the states are all zeroed, their covariance matrix is diagonalized (off diagonal terms are zeroed), and the diagonal elements are reinitialized. In the post-mission software, the changes in latitude, longitude and heading across the fix are estimated and used to correct the state vector. The covariance matrix is not changed. The changes in latitude, longitude and heading across the fix were computed using inertial data: when an INS position fix is detected, the post-mission software computed the discontinuity resulting from the fix by subtracting the value just before the fix from the value just after the fix. No information is required from the INS except the fix flag. However, using this procedure for computing heading resets produces erroneous results: the rates of change of heading due to wave action, etc. are too large to estimate small changes in heading errors. If accurate measurements, such as GPS, are available, the heading state estimates quickly return to their proper values. When less accurate measurements, such as Loran-C, are used, the incorrect adjustment of the heading state "permanently" disturbs the filter.

The post-mission position reset procedure was corrected by simply removing the heading reset (ie. the DIINS heading state is not modified in response to a MK29 position fix). The filtered position errors were very similar to the real-time results after this change was made.

Currently, the latitude, longitude and heading states computed by the SKOR filter are used to correct the corresponding DIINS states. Previously, it was thought that the SKOR state information was available via the DREO modifications to the MK29 digital outputs but not from standard Canadian Navy MK29s. It has now been confirmed that the state data is indeed available from all Canadian Navy systems. The deterministic corrections applied in DIINS in response to INS position fixes are more reliable when taken directly from the INS. This change did not alter the conclusions derived above.

It has thus far not been possible to reproduce the small latitude standard deviations found in the real-time results. Analysis showed that the estimated Time Difference (TD) error state standard deviations are approximately equal to the position error state standard deviations: $\sigma_{TDY} \simeq \sigma_{lat}$ and $\sigma_{TDZ} \simeq \sigma_{lon}$. A look at the Loran-C chart for the area of the sea trials, L/C 3001, shows that the hyperbolic line of position 5990–Y is at about 70 degrees azimuth and 5990–Z is at about 10 degrees. So, the correlation with latitude and longitude is reasonable.

There is nothing in the real-time results or the post-processed results to indicate why estimated latitude error state standard deviations are smaller than longitude:

1. Filter models for INS latitude and longitude errors are the same.
2. Measurement models for all TDs are the same.
3. The geometrical accuracy for latitude varies from slightly better to equal to the geometrical accuracy for longitude. Geometrical accuracy in this instance is defined as the projection of 10 millisecond TDs onto parallels of latitude and meridians of longitude, respectively. It is comprised of a gradient component related to ship's position versus the master-slave baseline and a second component related to the crossing angle of the two TDs being tracked and the orientation of the TDs relative to cardinal directions.
    a. The TD-Y gradient is slightly smaller than the TD-Z gradient, implying slightly better latitude accuracy.
    b. The crossing angle / orientation geometry implies slightly better longitude accuracy.

4. The DIINS Loran-C measurement covariance matrix is computed using receiver supplied Signal to Noise Ratios (SNRs) and a priori constant scale factors. All TDs have the same scale factor. TD measurement noise variances are computed by converting SNRs in decibels to unitless Noise to Signal Ratios (NSRs). The NSRs are rescaled by the a priori scale factors to derive range standard deviations. TD variances are computed by adding the Master range variance to each Slave range variance. Plots of the TD standard deviations show that the expected level of noise for TD-Y is very similar to TD-Z.
5. Residual standard deviations for TD-Y and TD-Z are approximately equal.

Fix of Anomalous Condition Since the post-processing software does not exhibit the anomalous behaviour under investigation, and since no way has been found to reproduce this behaviour, no fix can be recommended (if one is necessary). At present, the hardware required to test the real time system is not available. When possible, the real time system will be run with Loran-C to try to reproduce this condition.

4.7: Filter 7 (INS 2 with LC360 Loran-C Time Differences)

Recall that Filter 6 and Filter 7 differ only in the INS being filtered: the configuration of Filter 7 is exactly the same as Filter 6 (as described in the preceding section). In summary, 4 Loran-C TD measurements were activated and a total of 18 states were modelled.

All the anomalies reported above for Filter 6 also apply to Filter 7. No additional anomalies or insights were observed with Filter 7.

Chapter 5: REAL TIME FAILURE DETECTION, ISOLATION AND RECONFIGURATION

This chapter describes INS status messages available over the MK29 serial digital interface and their effects on ST95 real time results as well as the methods and results of the Failure Detection, Isolation and Reconfiguration (FDIR) task used in DIINS for ST95. Note that DIINS FDIR methods have been significantly changed, expanded and improved since ST95. A description of the updated DIINS FDIR processor can be found in reference [10].

The times in this chapter are given in our peculiar DIINS time format. Detailed conversion procedures were given in the Introduction. In brief, DIINS time is the time elapsed since the start of the year 1993, in seconds. It can be converted to Julian days as follows:

$$\text{JDay} = \frac{\text{DIINS time (in seconds)}}{86,400 \text{ (sec / day)}} (\text{Year} - 1993) \tag{31}$$

For example,

$$\text{JDay} = \frac{74,276,507}{86,400} (1995 - 1993) = 129.6818 \tag{32}$$

5.1: INS Status Messages

There are 14 status flags sent by the Sperry MK29 Mod 3 INSs. The precise meanings of some of them are not presently known. A request for information has been made to Sperry but no reply has yet been received. A description of each flag and the current best guess at their meanings are listed below. If there is any uncertainty in the interpretation of the flag, there will be a (??) symbol to indicate such.

1. Performance Degraded (??): According to the available literature, this flag is set to true when too much time has elapsed since the last position fix or when the SKOR estimate of the heading uncertainty exceeds 5 arc minutes. The time since last fix criterion is apparently not fixed but will range from 10–32 hours. Exactly how the criterion is set is unknown: it is likely a function

of the accuracy used for the last position fix. In addition, analysis of sea trial data has shown that the Performance Degraded flag may be set when neither of the known criteria are met. This may mean that limits given in the manuals are not the same as those currently implemented in the INSs. Or, perhaps there are other criteria by which this flag can set. Examples of the flag being set without meeting either of these criteria have been discussed in Chapters 4 and 3.

2. Fine Align Settled (??): The implications of this flag can be deduced from its name. No other information has been found. However, this flag is never set to false during Sea Trial 1995.
3. Inertial Well Settled (??): The meaning of this flag is not clear. It was set to FALSE (signalling a failure) several times during ST95. It is briefly discussed in Chapter 4.
4. System Mode: This flag has three possible values
   a. NON_NAVIGATE when the system is not ready to navigate (eg. the alignment procedure is not complete).
   b. GYRO_FINE_ALIGN (??) when the system in performing its fine alignment. Is this set when the system is acting as an gyrocompass (the alignment is complete but no position fixes have been supplied)?
   c. INERTIAL when the system is performing as an inertial navigator (the alignment is complete and position fixes have been supplied).
5. Velocity Damping Mode: The velocity damping loop can be activated and deactivated AUTOmatically using predefined algorithms or MANually Damped or Undamped by the operator.
6. Velocity Damping Enabled: This flag is self explanatory.
7. SKOR Mode: This flag has three possible values
   a. OFF when the SKOR filter is deactivated.
   b. STORE when the SKOR filter is activated but the results are stored in memory but the INS is not being reset.
   c. APPLY when the SKOR filter is activated and the INS is being reset with the SKOR results.
8. SKOR Status: This flag has four possible values
   a. PROCESS_INCOMPLETE is self-explanatory.

b. FIX_REJECT when a fix which has been input into the SKOR filter fails to meet expected criteria (residuals are larger than expected ??).
c. FIX_ACCEPT when a fix meets all the expected SKOR criteria.
d. UPDATE when the state covariances are being propagated between fixes.

9. Position Fix Source: This flag has five possible values

   a. MANUAL when the operator uses the display panel to input position fixes from an external source. This was source used in ST95.
   b. SATNAV when the Transit satellite system is used for position fixes via the SDI.
   c. OMEGA when the Omega system is used for position fixes via the SDI.
   d. LORAN when the Loran-C system is used for position fixes via the SDI.
   e. GPS when the GPS satellite system is used for position fixes via the SDI.

10. Speed Log Failure: This flag is self explanatory.
11. System Failure (??): This sounds like a fatal failure, but the criteria used are not known.
12. System On Internal Battery: This flag is set whenever the ship's power supply to the INS fails. The MK29 has an internal battery which is able to power essential functions for about 30 minutes.
13. Velocity Reference: This flag has three possible values

    a. MANUAL when the operator uses the display panel to input velocities from an external source. In Dockside mode, MANUAL velocity reference (zero velocity) is assumed.
    b. SPEED_LOG when speed log data is supplied over a synchro interface. This was the velocity reference used whenever the ship was sailing during ST95.
    c. SERIAL_DIGITAL_INTERFACE when velocities are supplied over the serial digital interface.

14. Dockside Mode: This flag is set by the operator via the display panel. The ship must be docked at a known position. In Dockside Mode, the INS will automatically perform position fixes every six minutes using the known position as reference.

Below is a list of all significant changes in INS status for most of Sea Trial 1995. Status flags for both INSs for all four runs have been analyzed. If any particular flag is not listed for a particular INS for a particular run, it can be assumed that the flag was initialized to the expected value and did not change over the period of the run. Summaries for both INSs for all runs are attached. Recall that all times are given in the DIINS time format described above.

5.1.1. Run 1, Days 129–132

MK29–1 (SN 617)

1. Manual SKOR position fixes were given at times
   a. 74,276,507 seconds — out of Dockside Mode;
   b. 74,286,275 seconds — Fix accepted,
   c. 74,293,751 seconds — Fix accepted,
   d. 74,300,702 seconds — Fix accepted,
   e. 74,308,416 seconds — Fix accepted,
   f. 74,315,373 seconds — Fix accepted,
   g. 74,344,615 seconds — Fix accepted,
   h. 74,451,837 seconds — Fix accepted,
2. The Performance Degraded flag is set to TRUE at times
   a. 74,281,735 seconds to 74,283,533 seconds — no known reason (time since Dockside is 1.45 hours, σ_{HDG} is 1 arcmin);
   b. 74,324,727 seconds to 74,326,526 seconds — no known reason (time since fix is 2.60 hours, σ_{HDG} is 1 arcmin);
   c. 74,418,594 seconds to 74,451,843 seconds — time since fix is 20.55 hours, σ_{HDG} is 11 arcmin (Why was the flag not set when σ_{HDG} exceeded 5 arcmin at about 74,386,000 seconds? A fix occurred at 74,451,837, causing the flag to be reset to FALSE.);
   d. 74,476,877 seconds to 74,478,672 seconds — no known reason (time since fix is 1.45 hours, σ_{HDG} is 3.5 arcmin);

3. The Inertial Well Settled flag is to FALSE set at times

   a. 74,329,032 seconds to 74,345,218 seconds;
   b. 74,356,727 seconds to 74,452,402 seconds;
   c. 74,454,201 seconds to the end of the run (approximately 74,491,097 seconds);

MK29–2 (SN 611)

1. Manual SKOR position fixes were given at times

   a. Approximately 74,276,000 seconds — out of Dockside Mode;
   b. 74,285,987 seconds — Fix REJECTED,
   c. 74,286,363 seconds — Fix accepted,
   d. 74,293,448 seconds — Fix accepted,
   e. 74,300,420 seconds — Fix accepted,
   f. 74,308,150 seconds — Fix accepted,
   g. 74,315,085 seconds — Fix accepted,
   h. 74,344,130 seconds — Fix accepted,
   i. 74,451,743 seconds — Fix accepted,

2. The Performance Degraded flag is to TRUE set at times

   a. 74,292,259 seconds to 74,296,004 seconds — no known reason (time since fix is 1.64 hours, σ_{HDG} is around 1 arcmin);
   b. 74,326,046 seconds to 74,327,844 seconds — no known reason (time since fix is 3.04 hours, σ_{HDG} is about 1.5 arcmin);
   c. 74,357,398 seconds to 74,359,197 seconds — no known reason (time since fix is 3.69 hours, σ_{HDG} is about 2 arcmin);
   d. 74,361,785 seconds to 74,363,583 seconds — no known reason (time since fix is 4.90 hours, σ_{HDG} is about 2 arcmin);
   e. 74,379,702 seconds to 74,383,368 seconds — no known reason (time since fix is 9.88 hours, σ_{HDG} is about 4 arcmin);
   f. 74,385,589 seconds to 74,387,734 seconds — time since fix is 11.52 hours, σ_{HDG} is about 5 arcmin;
   g. 74,389,956 seconds to 74,392,247 seconds — time since fix is 12.73 hours, σ_{HDG} is about 6 arcmin;
   h. 74,395,586 seconds to 74,451,753 seconds — time since fix is 14.29 hours, σ_{HDG} is from about 7 arcmin to about 18.5 arcmin;

i. 74,470,906 seconds to 74,472,625 seconds — time since fix is 5.30 hours, σ_{HDG} is about 3 arcmin;

There is no apparent reason for the recovery of the Performance Degraded flag between 74,385,000 seconds, when the SKOR heading standard deviation exceeds 5 arcmin) and the fix at 74,451,000 seconds.

3. The Inertial Well Settled flag is to FALSE set at times

   a. 74,328,445 seconds to 74,344,630 seconds;
   b. 74,356,140 seconds to 74,452,174 seconds;
   c. [illegible]453,613 seconds to the end of the run (approximately 74,491,097 seconds);

5.1.2. Run 2, Day 133

MK29–1 (SN 617)

1. Manual SKOR position fixes were given at times

   a. 74,622,714 seconds — out of Dockside Mode;
   b. 74,630,018 seconds — Fix accepted,
   c. 74,637,051 seconds — Fix accepted,
   d. 74,644,246 seconds — Fix accepted,
   e. 74,651,522 seconds — Fix accepted,
   f. 74,657,525 seconds — Fix accepted,

MK29–2 (SN 611)

1. Manual SKOR position fixes were given at times

   a. 74,622,598 seconds — out of Dockside Mode;
   b. 74,629,855 seconds — Fix accepted,
   c. 74,638,145 seconds — Fix accepted,
   d. 74,644,302 seconds — Fix accepted,
   e. 74,648,639 seconds — Fix accepted,
   f. 74,651,656 seconds — Fix accepted,
   g. 74,657,539 seconds — Fix accepted,

2. The Performance Degraded flag is to TRUE set at times
    a. 74,634,767 seconds to 74,638,945 seconds — no known reason (time since fix is 1.36 hours, σ_{HDG} is about 1 arcmin);
    b. 74,642,199 seconds to 74,643,997 seconds — no known reason (time since fix is 1.13 hours, σ_{HDG} is about 1 arcmin);
    c. 74,665,057 seconds to end of run (about 74,668,896 seconds) — no known reason (time since fix is 2.09 hours, σ_{HDG} is less than 2 arcmin.

5.1.3. Run 3, Days 134–135

MK29–1 (SN 617)

1. Manual SKOR position fixes were given at times
    a. 74,686,620 seconds — Fix accepted,
    b. 74,757,225 seconds — Fix accepted.
2. The SKOR σ_{HDG} exceeds 5 arcmin at about 74,728,000 seconds, recovers to about 1.5 arcmin after the fix at 74,757,225 seconds, then exceeds 5 arcmin again at about 74,795,000 seconds. However, the Performance Degraded flag stays at the FALSE setting though the run.
3. The Inertial Well Settled flag is to FALSE set at times
    a. Initialization (74,674,687) seconds to 74,687,273 seconds;
    b. 74,698,783 seconds to 74,757,771 seconds;
    c. 74,766,043 seconds to the end of the run (about 74,822,546 seconds);

MK29–2 (SN 611)

1. Manual SKOR position fixes were given at times
    a. 74,686,392 seconds — Fix accepted,
    b. 74,757,073 seconds — Fix accepted.
2. The Performance Degraded flag is to TRUE set at times
    a. 74,677,853 seconds to 74,679,652 seconds — no known reason (time since last **known** fix is 5.64 hours, σ_{HDG} is about 2.5 arcmin);
    b. The SKOR σ_{HDG} exceeds 5 arcmin at about 74,728,000 seconds, recovers to about 1.5 arcmin after the fix at 74,757,073 seconds, then exceeds 5

arcmin again at about 74,794,000 seconds. However, the Performance Degraded flag stays at the FALSE setting at all times after 74,679,652 seconds.

3. The Inertial Well Settled flag was to FALSE set at times
   a. Initialization (74,674,687) seconds to 74,687,045 seconds;
   b. 74,698,555 seconds to 74,757,542 seconds;
   c. 74,765,815 seconds to the end of the run (about 74,822,546 seconds);
4. The Velocity Damping Mode was
   a. Initialized to MANUAL (Damped) at 74,674,687 seconds;
   b. Changed to AUTO at time 74,689,953 seconds;
   c. Changed to MANUAL (UNdamped) at time 74,717,743 seconds;
   d. Changed to AUTO at time 74,717,745 seconds;
   e. Changed to MANUAL (Damped) at time 74,722,481 seconds;
   f. Changed to AUTO at time 74,724,087 seconds;
5. SKOR Mode was
   a. Initialized to APPLY at 74,674,687 seconds;
   b. Changed to OFF at time 74,717,743;
   c. Changed to APPLY at time 74,717,745;
6. The Velocity Reference was
   a. Initialized to MANUAL at 74,674,687 seconds;
   b. Changed to SPEED_LOG at time 74,717,745;

5.1.4. Run 4, Day 136

MK29–1 (SN 617)

1. The Performance Degraded flag is set to TRUE
   a. For the entire run, from initialization (74,830,965 seconds) to the end (about 74,881,067 seconds). There were no fixes during this run: the last known fix was at 74,757,225 seconds, 20.48 hours before the start of this run. The SKOR σ_{HDG} starts the run at about 11 arcmin and grows to about 25 arcmin by the end.

2. The Inertial Well Settled flag was to FALSE
   a. For the entire run, from initialization (74,830,965 seconds) to the end (about 74,881,067 seconds).
3. The SKOR Mode flag was
   a. Initialized to APPLY at 74,830,965 seconds;
   b. Changed to STORE at time 74,841,943 seconds;

MK29–2 (SN 611)

1. The Performance Degraded flag is to TRUE
   a. For the entire run, from initialization (74,830,964 seconds) to the end (about 74,881,067 seconds). There were no fixes during this run: the last known fix was at 74,757,073 seconds, 20.53 hours before the start of this run. The SKOR σ_{HDG} starts the run at about 11 arcmin and grows to about 25 arcmin by the end.
2. The Inertial Well Settled flag was to FALSE
   a. For the entire run, from initialization (74,830,964 seconds) to the end (about 74,881,067 seconds).
3. The Velocity Damping Mode was
   a. Initialized to MANUAL (Damped) at 74,840,711 seconds;
   b. Changed to AUTO at time 74,841,835 seconds;
4. The SKOR Mode flag was set to STORE at time 74,842,244 seconds;

5.2: Measurement Analysis (Hard Failures or Data Problems)

The measurements formed by the DIINS Kalman filters were used in two ways in the FDI processing: when there were no measurements formed using a particular sensor, data problems are assumed and a failure condition is declared; when a particular measurement's residual consistently failed its test, another failure condition was declared. Immediately following any change in sensor status, the filter classification algorithm was run to select the best filter for output. More detail is provided below.

5.2.1. Data Problems

When all of the criteria prescribed for the formation of a particular type of measurement (eg. INS / GPS positions) were met, counters for the INS sensor and the aiding sensor (eg. GPS) were incremented. Let's call these the MeasFormed counters. There is one such counter for each sensor in each filter. There are three conditions under which measurements were not formed:

1. When the data record for the INS or the aiding sensor had a bad status flag.
2. When the latest data record for either sensor had previously been used to form a measurement.
3. When there was a time mismatch between the INS and aiding sensor data records which exceeded the maximum allowed.

In the ST95 FDIR task, the TotalMeasFormed counters for all activated sensors were computed by summing the MeasFormed counters from each FDIR filter (ie. the INS1/INS2 filter, the INS1 unifilter and the INS2 unifilter). If the TotalMeasFormed counter for any sensor was zero and remained at zero for a prescribed length of time, the " . . . NO measurements being processed . . . " failure condition was declared for that sensor. For ST95, the time delay limit was set at 60 seconds. This failure condition was cancelled after the TotalMeasFormed counter exceeded zero for the same prescribed length of time

This condition was used in the filter classification algorithm but did NOT affect measurement processing in the filters.

The following is a summary of the FDIR output resulting from a sensor being rejected because of a lack of measurements:

1. Run 1: Julian Days 128–132,
   a. GPS — NO measurements being processed from 74,211,323 seconds to 74,334,135 seconds;
   b. GPS — NO measurements being processed from 74,334,432 seconds to 74,376,219 seconds;
2. Run 2: Julian Day 133,
   a. GPS — NO measurements being processed from 74,660,814 seconds to 74,660,904 seconds;
   b. GPS — NO measurements being processed from 74,661,806 seconds to 74,661,895 seconds;

3. Run 3: Julian Days 134–135,
    a. GPS — NO measurements being processed from 74,766,992 seconds to the end of the run (about 74,822,546 seconds);
    b. Loran-C — NO measurements being processed from 74,802,796 seconds to the end of the run (about 74,822,546 seconds);
4. Run 4: Julian Day 136,
    a. GPS — NO measurements being processed from 74,823,872 seconds to the end of the run (about 74,881,067 seconds) ;

The data problems with GPS were discussed in Chapter 2.

5.2.2. Hard Failures

For ST95, a DIINS hard sensor failure declaration required the following set of conditions:

1. For each particular measurement series (eg. INS1/GPS latitude), a measurement failure was declared after the series had failed its residual test X percent of the time over the last Y seconds. For ST95, the limits were set at 90% of the measurements over the last 5 minutes (except Filter 1).
2. Once a measurement series had been declared FAILED, the MeasFailed counters for the INS and aiding sensor (eg. INS1 and GPS) were incremented. At this point, the failed series would no longer be used to perform Kalman filter updates in its own filter (this did not affect other filters).
3. In the FDIR processor, the TotalMeasFailed counters for all activated sensors were computed by summing the MeasFailed counters from each FDIR filter (ie. the INS1/INS2 filter, the INS1 unifilter and the INS2 unifilter).
4. If the TotalMeasFormed and TotalMeasFailed counters were equal (but not zero), a "Hard Failure" condition was signalled. This indicated that every measurement type involving the failed sensor had failed in every filter (eg. at least one GPS position measurement had failed in every filter). At this point, all measurements in any filter using the failed sensor would be prevented from performing Kalman filter updates.

Even when measurements from a failed sensor were not used to update a Kalman filter, the residuals from those measurements were still computed and tested. The sensor's behaviour could be thus be monitored and, if appropriate, a

failed sensor could be reactivated. The reactivation procedure was very similar to that used to declare a failure:

1. For any particular measurement series (eg. INS1/GPS latitude) in filter (I), a measurement recovery was declared after the measurement had passed its residual test XX percent of the time over the last YY seconds. For ST95, the recovery limits were also set at 90% of the measurements over the last 5 minutes (except Filter 1).
2. After a series had RECOVERED in filter (I), the MeasFailed counters for the sensors involved (eg. INS1 and GPS) were NOT incremented. Thus, MeasFailed(I) < MeasFormed(I).
3. In the FDIR processor, the TotalMeasFailed and TotalMeasFormed counters for all activated sensors were computed by summing the MeasFailed and MeasFormed counters from each FDI filter.
4. Since MeasFailed(I) < MeasFormed(I), and since, in general, MeasFailed ≤ MeasFormed, the summations performed in the FDIR processor would necessarily produce the result TotalMeasFailed < TotalMeasFormed. Thus, the "Hard Failure" condition would be cancelled and the previously failed sensor could again be used for Kalman filter updates.

The following is a summary of the measurement failures for each filter during the real time DIINS processing during Sea Trail 1995:

1. Run 3: Julian Days 134–135,
    a. Filter 6 (INS1+LC):
        i. Loran-C — the Time Difference Measurement from Slave Y (TD-Y) failed from 74,675,161 seconds to 74,675,842 seconds.
2. Run 4: Julian Day 136,
    a. Filter 2 (INS1+PY+LC):
        i. Loran-C — TD-Z failed from 74,824,254 seconds to 74,825,209 seconds.
        ii. Loran-C — TD-Y failed from 74,825,608 seconds to 74,827,648 seconds.

b. Filter 6 (INS1+LC):

   i. Loran-C — TD-Z failed from 74,824,315 seconds to 74,869,659 seconds.
   ii. Loran-C — TD-Y failed from 74,831,755 seconds to 74,834,107 seconds.
   iii. Loran-C — TD-Y failed from 74,834,658 seconds to 74,836,329 seconds.
   iv. Loran-C — TD-Y failed from 74,836,844 seconds to the end of the run (about 74,881,067 seconds).

The following is a summary of the FDIR output regarding sensor hard failure isolation during the real time DIINS processing of Sea Trail 1995:

1. Run 4: Julian Day 136,

   a. INS1 — a Hard Failure Isolation followed by a Recovery is reported every few seconds from 74,824,259 seconds to 74,827,633 seconds;
   b. Loran-C — a Hard Failure Isolation followed by a Recovery is reported every few seconds from 74,824,256 seconds to 74,827,633 seconds;
   c. Both sensors have reportedly recovered at 74,827,635 seconds and all FDIR processes behave as expected until the end of the run — almost 15 hours later.

This rapid switching from hard fail to recovery to hard fail etc. described for Run 4 above should not be possible because of the time delay feature built into the FDIR processor. Analysis of the residual files indicates that

1. There is not so much as one failed residual for Filter 1 (INS1 + INS2); there are no measurement failures; and therefore, there should not be ANY INS1 or INS2 hard failures.
2. There are Loran-C measurement failures in Filters 2 and 6. However, there is not much correlation between the residual plots and FDIR messages.
3. Filter 2 summary:

   i. The TD-Z residuals alternate quickly between 0 and 500–1000 metres up until 74,825,200 seconds. This suggests that only a few Loran-C data messages are being used to form measurements. Perhaps there are timing problems similar to those identified for GPS. At the same time, the residual standard deviation grows from about 40 metres to about

1500 metres. After 74,825,200 seconds, the residual behaviour is more reasonable: the residuals look like noise with standard deviations of a few metres; the expected standard deviations range from about 6 to 20 metres.

ii. The TD-Y residuals alternate quickly between 0 and 700–500 metres from 74,825,200 seconds to 74,827,300 seconds. Again, this suggests that only a few Loran-C data messages are being used to form measurements. Over this period, the standard deviation grows from 50 metres to 600 metres. Up to 74,825,200 seconds, the residual behaviour is somewhat more reasonable. However, the residuals are erratic, varying from a few metres to several tens of metres. The standard deviations over this interval also range from 6 to 20 metres.

4. Filter 6 summary:

   i. TD-Y residuals seem to suffer from the same sparse measurement problem — they move quickly between ± 1000 metres. The standard deviations range from 6 to 60 metres.

   ii. TD-Z residuals all fail up until about 74,869,700 seconds. They grow up over 13 kilometres in size. Their standard deviations grow up to over 2 kilometres. After 74,869,700 seconds, residuals and standard deviations look good.

   Note that Filter 6 was never selected as a unifilter and thus should never have been used for hard failure detection.

The effects of this disturbing hard failure behaviour on the filter classification results are discussed below.

5.3: Chi-Squared Tests (Soft Failures)

DIINS executes a second level of FDIR testing which is designed to act on more subtle, slowly developing sensor errors which may not be detectable using the hard failure tests. The filters will often redistribute unmodelled errors to inappropriate Kalman states in situations where those unmodelled errors are slowly growing. To detect these "soft" sensor failures in DIINS, the state vector and its covariance matrix are compared with a "shadow filter" which uses exactly the same initial conditions and models as the primary filter. However, the shadow filter does not process any measurements. Without measurements, the shadow filter will propagate the initial conditions using the given models: it will not

be corrupted by faulty sensor information. The difference in state vectors and covariance matrices are used to compute a statistic which has a χ^2 distribution. The test is based on the premise

$$\Pr\left\{(\vec{x}-\vec{x}_s)^T(P-P_s)^{-1}(\vec{x}-\vec{x}_s) > \chi^2_{\alpha,m}\right\} = \alpha \tag{33}$$

where α is the significance level and m is the degrees of freedom (= number of states used). Thus, whenever the expression inside the probability brackets is true, it is assumed that one of the sensors in the filter has deviated from its initial model: the χ^2 test has failed. This test does not provide any information on the specific sensor that has failed. For ST95, the failed sensor was isolated in the FDIR task according to Table 18. where

Table 18 Soft Failure Isolation

F1-2	F1-1	F1	F0	F2	F2-1	F2-2	Soft Failure
X	-	X	X	0	0	0	INS1
0	0	0	X	X	-	X	INS2
X	O	X	O	X	O	X	A1
O	X	X	O	X	X	O	A2

1. F0 is the INS1 / INS2 filter;
2. F1 and F2 are the INS1 and INS2 unifilters, respectively;
3. FI-J is the filter based on INS(I) and using all aiding sensors except J.
4. A1 and A2 are the two aiding sensors.
5. O indicates that the filter passed the χ^2 test;
6. X indicates that the filter failed the χ^2 test;
7. — indicates that the results are not relevant to the isolation procedure.

Note that numerical problems are common in this test, especially at the start of the run. The matrix being inverted $(P - P_s)$ is the difference of two similar covariance matrices. At initialization, $(P - P_s) = 0$ and has no inverse. Shortly after initialization, $(P - P_s) \simeq 0$ resulting in bad conditioning, numerical problems and unreliable results. The DIINS software checks the conditioning of $(P - P_s)$ and prevents the inversion when the conditioning is poor. This will produce a "Chi-squared CONDITION failure".

The following is a summary of the χ^2 test failures (*soft failure detection*) for each filter during the real time DIINS processing during Sea Trial 1995:

1. Run 4: Julian Day 136,
   a. Filter 1 (INS1+INS2):
      i. χ^2 test failed from 74,823,813 seconds to 74,826,419 seconds.
      ii. χ^2 test failed from 74,826,422 seconds to 74,826,425 seconds.
   b. Filter 2 (INS1+PY+LC):
      i. χ^2 test failed from 74,823,849 seconds to 74,823,985 seconds.
      ii. χ^2 test failed from 74,824,064 seconds to 74,824,108 seconds.
      iii. χ^2 test failed from 74,824,126 seconds to 74,824,144 seconds.
      iv. χ^2 test failed from 74,824,172 seconds to 74,824,195 seconds.
      v. χ^2 test failed from 74,824,210 seconds to 74,824,290 seconds.
      vi. χ^2 test failed from 74,824,366 seconds to 74,824,369 seconds.
      vii. χ^2 test failed from 74,825,211 seconds to 74,826,686 seconds.
   c. Filter 3 (INS2+PY+LC):
      i. χ^2 test failed from 74,823,818 seconds to 74,826,724 seconds.
      ii. χ^2 test failed from 74,826,726 seconds to 74,826,729 seconds.
   d. Filter 5 (INS2+PY):
      i. χ^2 test failed from 74,823,897 seconds to 74,826,649 seconds.
   e. Filter 6 (INS1+LC):
      i. χ^2 test failed from 74,823,816 seconds to 74,829,328 seconds.
   f. Filter 7 (INS2+LC):
      i. χ^2 test failed from 74,823,816 seconds to 74,826,783 seconds.

The following is a summary of the FDIR output resulting from a sensor being rejected because of failed chi-squared tests (*soft failure isolation*) during the real time DIINS processing of ST95:

1. Run 4: Julian Day 136,
   a. INS2 — Soft Failure from 74,823,959 seconds (from χ^2 test failure in Filter 5 at 74,823,897 seconds) to 74,826,486 seconds (from χ^2 test recovery in Filter 1 at 74,826,425 seconds);

b. Loran-C —

i. Soft Failure from 74,823,911 seconds (from χ^2 test failure in Filter 2 at 74,823,849 seconds) to 74,824,046 seconds (from χ^2 test recovery in Filter 2 at 74,823,985 seconds);

ii. Soft Failure from 74,824,272 seconds (from χ^2 test failure in Filter 2 at 74,824,210 seconds) to 74,824,351 seconds (from χ^2 test recovery in Filter 2 at 74,824,290 seconds);

iii. Soft Failure from 74,825,273 seconds (from χ^2 test failure in Filter 2 at 74,825,211 seconds) to 74,826,747 seconds (from χ^2 test recovery in Filter 2 at 74,826,686 seconds);

5.4: Selection of the Best Filter for Output

The output filter selection process worked as expected until Run 4. For the first hour of Run 4, the sensor hard failure isolation went wild with INS1 and Loran-C both rapidly being flagged as failed and then recovered. This was discussed (but not explained) above. Coinciding with this failure isolation bizarre behaviour, the output filter selection also went wild — in a span of 3800 seconds, the selection changed 734 times. The hard failure isolation procedure seemed to be out of control at the start of Run 4 (Julian Day 136). To illustrate, the first few records in the Day 136 FDI log file will be examined in detail:

1. At time 74,823,872 seconds, the MEAS_PROCESSED flag for GPS was set to FALSE. The best available aiding sensor after this time was changed from GPS to Loran-C. GPS did not recover at any time during this run.
2. Loran-C experienced a soft failure at time 74,823,911 seconds. At this point, both aiding sensors have failed. With no healthy aiding sensor, there should have been no unifilters identified. But, for some unknown reason, the FDI log file indicates that the Filters 2 and 3 were identified as the unifilters, and the output filter was changed from 2 to 3. (The expected position accuracy for Filter 3 was much higher at this time.)
3. INS2 experienced a soft failure at time 74,823,959 seconds. At this point, only INS1 had not failed. Again, the filter classification chose Filters 2 and 3 as the unifilters even though there were no healthy aiding sensors. Even more perplexing, Filter 3 was chosen as the output filter even though it was based on the failed INS2.

4. Loran-C recovered from its soft failure at 74,824,046 seconds. Unifilters and output filter remained unchanged. Again, since INS2 had failed, Filter 2 should have been selected as the output filter.
5. Loran-C experienced its first hard failure at time 74,824,256 seconds. This corresponds with the failure of the TD-Z measurement in Filter 2. The filter classification behaved as expected in this case: with no healthy aiding sensors, the unifilters could not be identified; the output filter did not change (although it was still wrong).
6. INS1 experienced its first hard failure at time 74,824,259 seconds. All sensors had failed. No unifilters were selected. Output filter = 1 is wrong: it should have stayed at 3.
7. INS1 and Loran-C recovered from hard failures at time 74,824,261 seconds. Unifilters = 2 and 3 was OK. Output filter = 3 was wrong because it was based on failed INS2; filter 2 should have been picked.
8. INS1 and Loran-C both fail at time 74,824,264 seconds. Filter classification was OK: no unifilters were identified and the output filter stayed at 3.
9. The confusion continued.

None of the FDIR problems with Run 4 (Day 136) are reproducible, and they are not consistent with the software. The only conceivable cause is an inconsistency in the real-time system (eg. cross-compiler or optimizer) which is not present in the post-processing software. Several other bugs causing unexplained errors in other parts of DIINS have been found in the real-time environment. These bugs were only identified through careful, tedious debugging of the real-time system.

Since the 1995 sea trial, the DIINS FDIR algorithms and software have been completely redesigned. None of the unexpected results described above have been produced by the new system (but, of course, it could not be reproduced with the old system either). The new system has been extensively tested on recorded sea trial data and seems to operating flawlessly.

Chapter 6: SUMMARY

The above report has described several aspects of the DIINS sea trial which was run from May 2–15, 1995.

The differential GPS system which served as the position reference was described and analyzed. In summary,

1. The system was owned and operated by the University of Calgary. Final positions were computed post-mission and delivered to DREO at the end of August 1995.
2. The system consisted of two 3951 NovAtel GPSCard™ receivers: one shore based set near the harbour at Esquimalt, the second on board the research ship, Endeavour.
3. Inter-receiver baselines ranged from a few kilometres to 250 kilometres.
4. The differential GPS reference was available on average about 78% of the time.
5. Based on analysis of ST95 residuals and analysis performed by the U of C using a carrier phase based GPS reference accurate to about 0.1 metre, the differential GPS reference for ST95 is estimated to be accurate at 0.5 to 1.0 metre level out to about 200 kilometres.

All three DIINS navigation sensors were described and their performance was analyzed.

1. The Sperry MK29 Mod 3 Ship's Inertial Navigation System:
    a. The operator should note that:
        i. The MK29s should NOT be held in Manual Damped mode through any kind of manoeuvre.
        ii. The SKOR position standard deviations may be too optimistic if the system has been running for an extended period of time in an undamped state — the SKOR filter seems to propagate them under the assumption that velocity damping is enabled.

b. Position Fixes:

   i. Fixes were entered at rates ranging from 2 to 34 hours.

   ii. All fix coordinates were given standard deviations of 0.3 nautical miles (550 metres). Fix errors (comprised of reference errors, transcription delays and quantization errors) were typically about 100 metres. No bad fixes were accepted during ST95 (although they have been found in earlier sea trials). A conservative value for the standard deviations was chosen to minimize the risk of corrupting the INSs with a bad fix.

   iii. Fix reliability (defined as the degree to which the post-fix position errors are represented by the SKOR standard deviation estimates) was found to be largely a function of pre-fix reliability: the better the pre-fix reliability, the better the post-fix reliability. The size of the pre-fix errors had little impact on the post-fix reliability provided the pre-fix SKOR variances were representative of the pre-fix errors.

   iv. The operator will often be able to estimate the fix residual using data from the reference sensor and SKOR data from the MK29 display. The estimated residual can be used to decide whether or not a rejected automatic fix should be accepted. The estimated residual can be computed before entering a manual fix. It can then be used to check the health of the reference sensor or to carefully adjust σ_{fix} to improve post-fix SKOR reliability.

   v. The rule-of-thumb fix rejection test used during ST95 is ineffective except in the unlikely event of the automatic rejection manual fix (this will only happen if the fix coordinates are given in grossly in error — in the wrong hemisphere, for example, or possibly, if an unreasonably small fix sigma is input).

c. Status flags: two status flags output over the MK29s' SDI were set to failure mode.

   i. The *Performance Degraded* flag is set in response to circumstances other than those described in Sperry manuals: according to the manuals, it is set if a fix has not been entered for some long period of time (eg. 10 hours) or the SKOR heading state standard deviation exceeds 5 arc minutes. However, the during ST95, it was set on

several occasions when neither of these conditions was met. There may some connection with the velocity damping status.

ii. The *Inertial Well Settled* flag is changed to FALSE several times during ST95. No documentation has been found which explains the criteria used to set this flag. There are no obvious correlations between the setting of this flag and INS performance.

d. DIINS' models: since the INSs are the primary sensor for DIINS' filters, INS position, velocity and attitude errors are defined as the system states. MK29 accelerometer and gyro biases were modelled as first order Gauss-Markov processes as follows:

i. X and Y accelerometer biases — standard deviation = 100 microg's, correlation time = 10 hours.

ii. X and Y gyro biases — standard deviation = 0.001 degrees per hour, correlation time = 24 hours.

iii. Z gyro bias — standard deviation = 0.003 degrees per hour, correlation time = 24 hours.

2. The Trimble TASMAN P(Y) code GPS receiver:

a. Position errors were computed relative to the dGPS reference. RMS differences based on all sea trial data were calculated as 5 metres horizontally and 7–8 metres vertically.

b. Position standard deviation estimates computed by the receiver agree well the dGPS position differences.

c. In the DIINS filters, this receiver's position errors are modelled as a first order Gauss-Markov process. In ST95 a standard deviation of 7 metres and correlation time of 120 seconds were used; measurement noise was set at 2 metres. This model was adjusted in DIINS' software through the use of the receiver's own accuracy estimates. This process was not operating properly during ST95 because of software bugs (which have since been discovered and corrected).

3. Internav LC360 Loran-C receiver:

a. Time Differences are output from the LC360 for processing in DIINS.

b. Errors in TDs are dominated by systematic (ASF) errors which are position dependent. TD 5990–Y is much more susceptible to large systematic

errors because its signal travels a longer distance over land. Better modelling of ASFs or some form of ASF calibration could dramatically reduce the size of TD errors.

c. DIINS has 4 TD error states, one for each slave in any particular chain. Off the south west coast of Vancouver Island, only 5990–Y and 5990–Z TDs are available.

d. For ST95, each TD error was modelled as a first order Gauss-Markov process with a standard deviation of 185 metres and a correlation time of 2.5 hours. Since the errors being modelled are position dependent, the correlations should be position dependent. However, that would be difficult to implement in DIINS. The TD measurements were assigned standard deviations of 10 metres. This model was borrowed form MINS. New models have been derived from DIINS sea trial data. But, since all trials were carried out in the same area, the new models developed are too site specific and not acceptable in other parts of the world.

The real time DIINS filters generally performed well. The biggest problem with real time processing occurred when the Data Collection Task for the Trimble GPS receiver flagged good data as stale, thus preventing the filters from using it. Subsequent analysis has shown that the GPS data staleness limit was too short and it has now been lengthened. Bugs were also found in the software procedure which modifies the GPS error model according to position variances transmitted from the receiver. This problem has also been fixed.

The two filters which used Loran-C as the only aiding sensor (Filters 6 and 7) exhibited unusual error behaviour. No explanation for the behaviour could be found. Nor could it be reproduced post-mission. It is recommended that these filters be thoroughly tested using the DIINS real time system to try to reproduce and repair the problem.

The failure detection processes worked well:

1. There were several Loran-C TD measurement failures ranging from a few seconds duration to 12 hours.
2. There were several soft failure detection failures ranging from a few seconds duration to 1.5 hours.

The soft failure isolation algorithms identified sensor failures in a manner consistent with the failure detection results: INS 2 was flagged for about 40 minutes during Run 4; Loran-C was flagged three times during Run 4, twice for

2–3 minutes and once for about 25 minutes. However, the hard failure isolation process went wild for about an hour during Run 4: the INS1 and Loran-C hard failure isolation messages were flashed every few seconds over this span of time. These messages seem to have no correlation to hard failure detection results and do not behave in a manner defined in the software — there is no explanation for this behaviour. Filters 2 and 6 do suffer from some strange Loran-C residual problems at this time that is probably connected to the spurious hard failure isolation messages but the connection is not clear. Again, post-mission processing did not produce this strange output. This has also been attributed to unidentified problems in the real time system. The good news is that the entire FDIR system has been rewritten since ST95 so the chances of this behaviour being repeated are slim.

Identification of the best filter for output worked as expected except during the anomalies with INS1 and Loran-C hard failure isolation just described: during this period of instability, the best filter selection messages were also changing unpredictably.

References

[1] D. A. G. Arden, "KALMAN FILTER DESIGN OF THE DUAL INERTIAL INTEGRATED NAVIGATION SYSTEM," Tech. Rep. Report No. DAC-02–97, Dale Arden Consulting, 1997. Prepared for the Defence Research Establishment Ottawa.

[2] D. A. G. Arden, "DIINS FAILURE DETECTION ISOLATION AND RECONFIGURATION PROCEDURES," Tech. Rep. Report No. DAC-03–97, Dale Arden Consulting, 1997. Prepared for the Defence Research Establishment Ottawa.

[3] G. Lachapelle and M. E. Cannon, "Field Report on Differential GPS Position and Attitude for Sea Trial Evaluation of DIINS (1995)," tech. rep., Department of Geomatics Engineering, University of Calgary, May 25 1995.

[4] H. Sun, G. Lachapelle, and M. E. Cannon, "Final Report on Differential GPS Position and Attitude for Sea Trial Evaluation of DIINS (1994)," tech. rep., Department of Geomatics Engineering, University of Calgary, Mar. 6 1995.

[5] G. Lachapelle and M. E. Cannon, "Final Report on Differential GPS Position and Attitude for Sea Trial Evaluation of DIINS (1995)," tech. rep., Department of Geomatics Engineering, University of Calgary, Aug. 26 1995.

[6] Sperry Marine Inc., Charlottesville, VA 22906, *MK-29 Mod 1S Advanced Stabilized Gyrocompass System and MK-29 Mod 3 Inertial Navigation System Operation and Service Manual*, ja17–4324b ed., Jan. 1988.

[7] D. DiFilippo, "DIINS Memo #4: Mark 29 Mod 3 Design Features Revisited." Unpublished, Mar. 1992.

[8] Sperry Marine Systems, *Proposal for Low Cost Ship's Inertial MK29 Mod 2*, ja18500–4035 ed., Feb. 1981.

[9] Sperry Marine Inc., Charlottesville, VA 22901, *MK29 MOD 3 Inertial Operation Guide*, May 1992.

[10] D. A. G. Arden, "Analysis of Methods Used in DIINS to Correct INS Positions." Unpublished, Jan. 1996.

Appendix A Post-Mission Position Fix Data Summaries

This appendix contains summaries of all ST95 position fixes. The following data summaries are copies of ASCII files "/home/Moe/DIINS/SEA-TRIALS/MAY95/MK29-1/DMK29-1_ST95-all.FixSummary" and "/home/Moe/DIINS/SEA-TRIALS/MAY95/MK29-2/DMK29-2_ST95-all.FixSummary" which were both created using program "/home/Moe/DIINS/SEA-TRIALS/MAY95/MK29-1/FixSummary.ada". The original format was corrupted when the files were read into Publisher™ — no effort was expended to improve the format here.

The data entries, in order, are

1. *Time* — DIINS time in seconds.
2. *Lat(Lon)_Err* — the error in INS latitude (longitude) at *Time*. This is approximately equal to the SKOR residual (only in the first line of each couplet).
3. *SKOR_Lat(Lon)* — the SKOR estimate of the latitude (longitude) error state at *Time*.
4. *SKOR_SD* — the SKOR estimate of the latitude (longitude) standard deviation at Time.
5. *Lat(Lon)_Test* — the estimated residual test statistic

$$\hat{m} = \frac{\hat{r}}{\hat{\sigma}_r} \simeq \frac{\hat{r}}{\sqrt{\sigma^2(-) + \sigma^2_{fix}}} \tag{34}$$

   where $\hat{r}$ is approximated by *Lat(Lon)_Err*, $\sigma(-)$ is the SKOR coordinate standard deviation immediately before the fix was entered (*SKOR_SD* on the first line of each couplet), and $\sigma_{fix} = 550$ metres.
6. *Rad_Err* — is the radial position error computed from *Lat_Err* and *Lon_Err*.

Notice that there are two entries for every fix. The first was computed when the fix was entered into the MK29; the second when the fix was accepted. The *Lat(Lon)_Test* statistics were not computed in the second line of each couplet because they require $\sigma(-)$. In the second lines, entries *Lon_Err*, *SKOR_Lon* and *SKOR_SD*—Lon are shifted once to the left, *Rad_Err* is shifted twice.

There were no position fixes entered during Run 4, Day 136.

A.1: MK29–1 Fix Summary

Time Lat_Err SKOR_Lat SKOR_SD Lat_Test Lon_Err SKOR_Lon SKOR_SD Lon_Test Rad_Err

Run 1, Days 129–132

74286275.5 1584.5 0.9 613.1 1.9 41.0 3.6 255.9 0.1 1585.0
74286281.9 743.0 840.1 413.3 -163.8 204.4 215.2 760.8
74293751.9 928.8 840.1 851.8 0.9 186.8 204.4 425.0 0.3 947.4
74293763.5 280.0 652.0 456.7 -41.2 230.6 298.4 283.0
74300702.3 262.7 652.0 862.7 0.3 -295.1 230.6 488.3 -0.4 395.1
74300707.5 153.4 109.4 461.2 -229.7 -66.0 340.9 276.2
74308416.9 -1865.3 109.4 917.0 -1.7 -890.8 -66.0 520.9 -1.2 2067.1
74308422.0 -403.3 -1462.2 471.1 -278.8 -613.1 359.9 490.3
74315373.7 -1744.5 -1462.2 909.7 -1.6 -949.8 -613.1 523.6 -1.2 1986.3
74315378.8 -349.6 -1396.2 471.1 -326.2 -623.1 362.6 478.1
74344615.3 -1169.6 -1396.2 3223.8 -0.4 -4319.7 -623.1 2696.6 -1.6 4475.3
74344621.7 458.4 -1626.8 525.4 -553.1 -3765.5 511.8 718.4
74451837.0 884.3 -1626.8 6721.6 0.1 -5143.4 -3765.5 20693.9 -0.2 5218.9
74451842.1 45.1 835.6 552.5 -29.4 -5113.8 558.9 53.9
74490938.5 -5457.6 835.6 5252.2 -1.0 -8406.9 -5113.8 6548.9 -1.3 10023.0
74490942.3 -10.5 -5442.1 260.4 -11.8 -8382.8 260.4 15.8

Run 2, Day 133

74630018.1 510.7 5.4 413.3 0.7 232.4 -3.6 177.2 0.4 561.1
74630030.9 305.0 203.5 340.9 173.6 62.4 167.3 351.0
74637051.7 799.8 203.5 784.0 0.8 757.6 62.4 356.3 1.1 1101.7
74637055.5 167.0 632.1 447.6 451.6 307.5 264.1 481.4
74644246.6 -574.1 632.1 860.9 -0.6 -143.3 307.5 463.9 -0.2 591.7
74644253.0 -151.4 -425.9 459.4 -10.0 -136.5 324.6 151.8
74651522.1 -1564.6 -425.9 873.6 -1.5 -896.1 -136.5 502.8 -1.2 1803.0
74651528.5 -397.1 -1166.5 463.9 -336.6 -558.0 351.8 520.6

74657525.3 -994.0 -1166.5 814.8 -1.0 -1081.7 -558.0 480.2 -1.5 1469.1
74657530.4 -147.7 -847.3 457.6 -526.2 -554.3 350.0 546.5

Run 3, Days 134–135

74686621.0 2559.8 -847.3 3203.0 0.8 633.7 -554.3 2677.6 0.2 2637.0
74686624.8 220.7 2339.4 523.6 -213.2 848.2 510.0 306.9
74757225.7 -125.0 2339.4 5018.8 0.0 -6967.8 848.2 11867.1 -0.6 6969.0
74757233.4 83.3 -207.1 551.6 -50.5 -6915.2 558.0 97.5

A.2: MK29–2 Fix Summary

Time Lat_Err SKOR_Lat SKOR_SD Lat_Test Lon_Err SKOR_Lon SKOR_SD Lon_Test Rad_Err

Run 1, Days 129–132

74286363.8 162.9 -4.5 584.2 0.2 119.5 -0.9 241.4 0.2 202.0
74286488.0 31.8 90.4 413.3 142.4 17.2 212.5 145.9
74293448.6 93.8 90.4 851.8 0.1 1461.0 17.2 423.2 2.1 1464.0
74293481.9 -230.2 329.2 456.7 1012.5 443.1 296.6 1038.4
74300420.8 -514.3 329.2 862.7 -0.5 -23.4 443.1 487.4 0.0 514.8
74300424.6 -105.0 -411.5 460.3 68.0 -87.7 340.0 125.1
74308150.7 -449.1 -411.5 917.0 -0.4 -223.0 -87.7 520.9 -0.3 501.4
74308154.5 -78.5 -371.7 471.1 -87.8 -134.7 359.9 117.7
74315085.7 167.9 -371.7 880.8 0.2 -142.2 -134.7 510.9 -0.2 220.0
74315089.6 101.6 66.0 467.5 -122.1 -21.7 359.9 158.8
74344130.2 107.0 66.0 3224.7 0.0 -15.8 -21.7 2694.8 0.0 108.1
74344135.3 109.9 -4.5 525.4 -3.4 -14.5 510.9 110.0
74451743.5 -361.8 -4.5 6742.4 -0.1 5611.9 -14.5 20797.0 0.3 5623.5
74451752.4 -49.3 -310.2 552.5 -11.8 5618.4 558.9 50.7

Run 2, Day 133

74629855.3 890.5 -4.5 413.3 1.3 547.2 0.0 177.2 0.9 1045.2
74629859.1 522.7 365.3 340.9 434.4 119.4 167.3 679.6
74638145.8 74.0 365.3 901.6 0.1 34.9 119.4 416.0 0.1 81.8
74638150.9 34.1 39.8 463.9 7.3 28.9 284.9 34.9
74648639.3 -579.1 39.8 1134.0 -0.5 -1034.0 28.9 631.2 -1.2 1185.1
74648643.1 -25.3 -553.4 485.6 -520.1 -513.6 371.7 520.7
74651656.2 331.6 -553.4 607.7 0.4 -697.7 -513.6 421.4 -1.0 772.5
74651660.0 211.1 124.8 412.4 -506.3 -189.0 331.9 548.5
74657539.1 192.2 124.8 731.6 0.2 315.2 -189.0 432.3 0.4 369.2
74657542.9 117.5 73.2 444.0 200.9 116.7 332.8 232.8

Run 3, Day 134–135

74686392.8 1353.1 73.2 3214.8 0.4 -703.3 116.7 2702.0 -0.3 1525.0
74686396.6 203.7 1148.5 523.6 -306.3 -400.6 510.0 367.8
74757073.0 4304.1 1148.5 5020.7 0.9 -7008.0 -400.6 11864.4 -0.6 8224.2
74757079.5 207.7 4093.8 551.6 -76.5 -6926.0 558.0 221.4

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Dale Arden Consulting
Suite 200, 300 March Road
Kanata ON K2K 2E2

2. SECURITY CLASSIFICATION (overall security classification of the document, including special warning terms if applicable)

UNCLASSIFIED

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C or U) in parentheses after the title.)

Analysis of Real Time Results from DIINS 1995 Sea Trials (U)

4. AUTHORS (Last name, first name, middle initial)

Arden, Dale A. G.

5. DATE OF PUBLICATION (month and year of publication of document)

August 1997

6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.)

110

6b. NO. OF REFS (total cited in document)

10

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Final Technical Report

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.)

Defence Research Establishment Ottawa
Ottawa ON K1A 0Z4

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant)

1bd11

9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)

W7714-5-9909

10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DAC-01-97 DRED CONTRACTOR REPORT 1999-116

10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Distribution limited to defence departments and defence contractors; further distribution only as approved
( ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
( ) Distribution limited to government departments and agencies; further distribution only as approved
( ) Distribution limited to defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)

UNLIMITED

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

The Dual Inertial Integrated Navigation System (DIINS) is a Kalman filter based, multi-sensor, optimally integrated navigation system developed at the Defence Research Establishment Ottawa. DIINS has been designed to provide a single source of navigation information to other ship board systems with minimum data latency to improve navigation reliability and accuracy, and to provide automatic sensor failure detection.

This report describes the DIINS sea trials which covered the time period from May 2, 1995 to May 15, 1995. Onboard the ship for these trials were two MK-29 MOD 3 INS's built by Sperry Marine Inc., a Trimble Navigation P(Y) code military GPS receiver and an Internav LC360 Loran-C receiver. Also onboard was a differential GPS system operated by the University of Calgary. This sensor acted as the position reference system and was used as a DIINS aiding sensor in post-mission tests.

This report begins with a description of all the navigation sensors involved in the trial. This is followed by a description of the real-time DIINS filter results. The report concludes with an analysis of the real-time FDIR results.

#512201

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Navigation, Kalman Filtering, Inertial Navigation, GPS, Loran-C, error analysis, Failure Detection, Failure isolation, FDIR, Sea trials